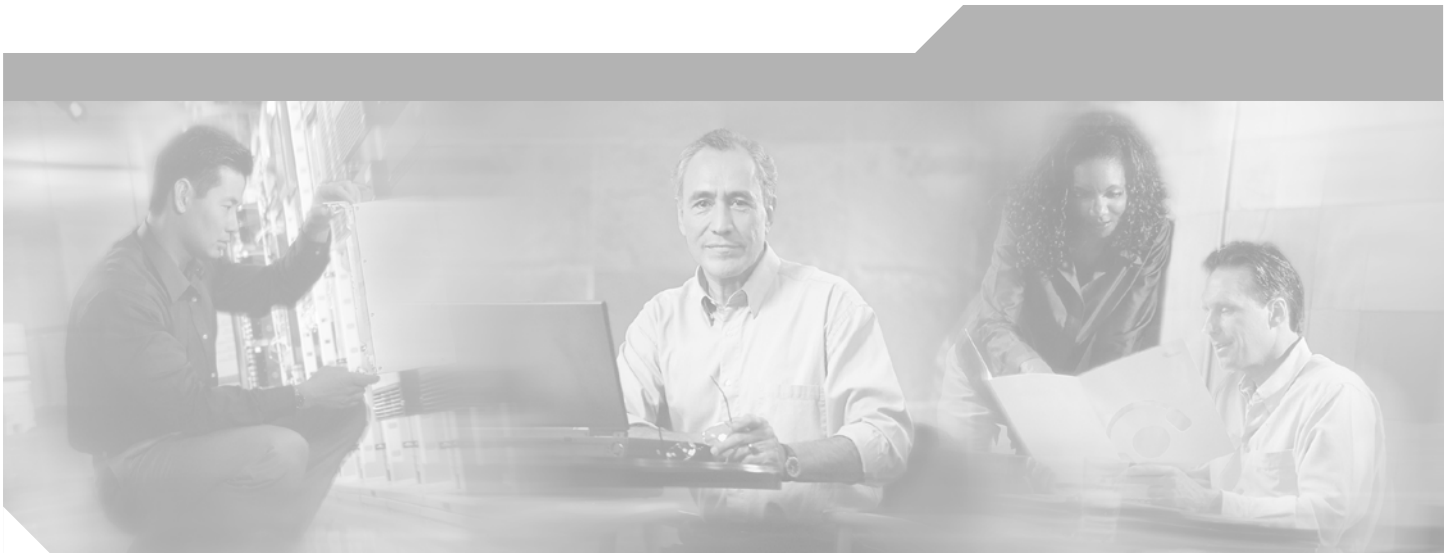

Data Center Networking: Integrating Security, Load Balancing, and SSL Services Using Service Modules

Solutions Reference Network Design
March, 2003

Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

Customer Order Number: 956639

VISIT…

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

CCIP, the Cisco Arrow logo, the Cisco *Powered* Network mark, the Cisco Systems Verified logo, Cisco Unity, Follow Me Browsing, FormShare, iQ Breakthrough, iQ Expertise, iQ FastTrack, the iQ Logo, iQ Net Readiness Scorecard, Networking Academy, ScriptShare, SMARTnet, TransPath, and Voice LAN are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, Discover All That's Possible, The Fastest Way to Increase Your Internet Quotient, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, LightStream, MGX, MICA, the Networkers logo, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, SlideCast, StrataView Plus, Stratm, SwitchProbe, TeleRouter, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0208R)

Data Center Networking: Integrating Security, Load Balancing, and SSL Services Using Service Modules
Copyright © 2003, Cisco Systems, Inc.
All rights reserved.

Preface i

Target Audience i

Document Organization i

Obtaining Documentation i

World Wide Web ii

Documentation CD-ROM ii

Ordering Documentation ii

Documentation Feedback ii

Obtaining Technical Assistance iii

Cisco.com iii

Technical Assistance Center iii

Cisco TAC Web Site iv

Cisco TAC Escalation Center iv

CHAPTER 1 **Data Center Overview — Integrating Security, Load Balancing, and SSL Services using Service Modules 1-1**

Benefits of Building Data Centers 1-1

Data Centers in the Enterprise 1-2

Data Center Architecture 1-3

Aggregation Layer 1-6

Front-End Layer 1-7

Application Layer 1-7

Back-End Layer 1-8

Storage Layer 1-8

Metro Transport Layer 1-9

Distributed Data Centers 1-9

Data Center Services 1-10

Infrastructure Services 1-10

Metro Services 1-10

Layer 2 Services 1-10

Layer 3 Services 1-11

Intelligent Network Services 1-11

Application Optimization Services 1-11

Storage Services 1-12

Security Services 1-12
Management Services 1-14
Summary 1-14

CHAPTER 2 **Integrating the Firewall Service Module 2-1**

Terminology 2-1
Overview 2-1
Deployment Scenarios 2-2
FWSM - MSFC Placement 2-4
MSFC-Outside 2-4
MSFC-Inside 2-5
FWSM - CSM Placement 2-5
Redundancy 2-6
Configurations Description 2-7
Common Configurations: Layer 2/Layer 3 2-7
Configuring VLANs 2-7
Configuring Trunks 2-8
Configuring IP Addresses 2-8
Configuring Routing 2-8
Configuring NAT 2-9
Configuring Redundancy 2-10
Intranet Data Center - One Security Domain 2-11
Internet Edge Deployment - MSFC-Inside 2-12
Multiple Security Domains / Multiple DMZs 2-12
Configurations 2-14
Intranet Data Center - One Security Domain 2-14
Aggregation1 2-15
Aggregation2 2-18
FWSM1 2-20
FWSM2 2-21
Internet Edge Deployment - MSFC Inside 2-22
Aggregation1 2-22
Aggregation2 2-25
FWSM1 2-27
FWSM2 2-28
Multiple Security Domains - Shared Load Balancer 2-29
Aggregation1 2-29
Aggregation2 2-32
FWSM2 2-36

CHAPTER 3 **Integrating the Content Switching Module 3-1**

Overview 3-1
What is the CSM 3-1
CSM Requirements 3-1
Interoperability Details 3-2
Data Center Network Infrastructure 3-2
Content Switching Interoperability Goals 3-3
Transparency 3-3
Scalability 3-3
High Availability 3-3
Performance 3-4
How the MSFC Communicates with the CSM 3-4
CSM Deployment 3-5
Aggregation Switches 3-5
Deployment Modes 3-6
Bridge Mode 3-6
Secure Router Mode 3-7
One Arm Mode 3-8
Server CSM MSFC Communication 3-8
High Availability 3-9
NAT (Network Address Translation) 3-10
Recommendations 3-10
CSM High Availability 3-11
Multi-Tier Server Farm Integration 3-13

CHAPTER 4 **Integrating the Content Switching and SSL Services Modules 4-1**

Terminology 4-1
Overview 4-1
Traffic Path 4-2
CSM SSL Communication 4-3
SSL MSFC communication 4-3
SERVERS CSM MSFC Communication 4-4
Redundancy 4-5
Security 4-6
Scalability 4-6
Data Center Configurations Description 4-7
Topology 4-7
Layer 2 4-9
Configuring VLANs on the 6500 4-10

Configuring VLANs on the CSM 4-11
Configuring VLANs on the SSLSM 4-11
Layer 3 4-12
Configuring IP Addresses on the MSFCs 4-12
Configuring IP Addresses on the CSM 4-12
Configuring IP Addresses on the SSLSM 4-12
Layer 4 and 5 4-12
CSM Configuration to Intercept HTTPS Traffic 4-13
SSLSM Configuration 4-13
Load Balancing the Decrypted Traffic 4-13
Returning Decrypted HTTP Responses to the SSLSM 4-14
Security 4-14
Multiple VIPs 4-15
Persistence 4-16
Configurations 4-16
Aggregation1 4-17
Aggregation2 4-21
SSL Offloader 1 4-25
SSL Offloader 2 4-25

INDEX

Preface

This Solution Reference Network Design (SRND) provides a description of the design issues related to integrating service modules in the data center.

Target Audience

This publication provides solution guidelines for enterprises implementing Data Centers with Cisco devices. The intended audiences for this design guide include network architects, network managers, and others concerned with the implementation of secure Data Center solutions, including:

- Cisco sales and support engineers
- Cisco partners
- Cisco customers

Document Organization

This document contains the following chapters:

Chapter or Appendix	Description
Chapter 1, “Data Center Overview — Integrating Security, Load Balancing, and SSL Services using Service Modules”	Provides an overview of data centers.
Chapter 2, “Integrating the Firewall Service Module”	Provides deployment recommendations for the Firewall Service Module (FWSM).
Chapter 3, “Integrating the Content Switching Module”	Provides deployment recommendations for the Content Switching Module (CSM).
Chapter 4, “Integrating the Content Switching and SSL Services Modules”	Provides deployment recommendations for the SSL Service Module (SSLSM).
Appendix A, “SSLSM Configurations”	SSLSM Configurations

Obtaining Documentation

These sections explain how to obtain documentation from Cisco Systems.

World Wide Web

You can access the most current Cisco documentation on the World Wide Web at this URL:

http://www.cisco.com

Translated documentation is available at this URL:

http://www.cisco.com/public/countries_languages.shtml

Documentation CD-ROM

Cisco documentation and additional literature are available in a Cisco Documentation CD-ROM package, which is shipped with your product. The Documentation CD-ROM is updated monthly and may be more current than printed documentation. The CD-ROM package is available as a single unit or through an annual subscription.

Ordering Documentation

You can order Cisco documentation in these ways:

- Registered Cisco.com users (Cisco direct customers) can order Cisco product documentation from the Networking Products MarketPlace:

  http://www.cisco.com/cgi-bin/order/order_root.pl

- Registered Cisco.com users can order the Documentation CD-ROM through the online Subscription Store:

  http://www.cisco.com/go/subscription

- Nonregistered Cisco.com users can order documentation through a local account representative by calling Cisco Systems Corporate Headquarters (California, U.S.A.) at 408 526-7208 or, elsewhere in North America, by calling 800 553-NETS (6387).

Documentation Feedback

You can submit comments electronically on Cisco.com. In the Cisco Documentation home page, click the **Fax** or **Email** option in the "Leave Feedback" section at the bottom of the page.

You can e-mail your comments to bug-doc@cisco.com.

You can submit your comments by mail by using the response card behind the front cover of your document or by writing to the following address:

Cisco Systems
Attn: Document Resource Connection
170 West Tasman Drive
San Jose, CA 95134-9883

We appreciate your comments.

Obtaining Technical Assistance

Cisco provides Cisco.com as a starting point for all technical assistance. Customers and partners can obtain online documentation, troubleshooting tips, and sample configurations from online tools by using the Cisco Technical Assistance Center (TAC) Web Site. Cisco.com registered users have complete access to the technical support resources on the Cisco TAC Web Site.

Cisco.com

Cisco.com is the foundation of a suite of interactive, networked services that provides immediate, open access to Cisco information, networking solutions, services, programs, and resources at any time, from anywhere in the world.

Cisco.com is a highly integrated Internet application and a powerful, easy-to-use tool that provides a broad range of features and services to help you with these tasks:

- Streamline business processes and improve productivity
- Resolve technical issues with online support
- Download and test software packages
- Order Cisco learning materials and merchandise
- Register for online skill assessment, training, and certification programs

If you want to obtain customized information and service, you can self-register on Cisco.com. To access Cisco.com, go to this URL:

http://www.cisco.com

Technical Assistance Center

The Cisco Technical Assistance Center (TAC) is available to all customers who need technical assistance with a Cisco product, technology, or solution. Two levels of support are available: the Cisco TAC Web Site and the Cisco TAC Escalation Center.

Cisco TAC inquiries are categorized according to the urgency of the issue:

- Priority level 4 (P4)—You need information or assistance concerning Cisco product capabilities, product installation, or basic product configuration.
- Priority level 3 (P3)—Your network performance is degraded. Network functionality is noticeably impaired, but most business operations continue.
- Priority level 2 (P2)—Your production network is severely degraded, affecting significant aspects of business operations. No workaround is available.
- Priority level 1 (P1)—Your production network is down, and a critical impact to business operations will occur if service is not restored quickly. No workaround is available.

The Cisco TAC resource that you choose is based on the priority of the problem and the conditions of service contracts, when applicable.

Cisco TAC Web Site

You can use the Cisco TAC Web Site to resolve P3 and P4 issues yourself, saving both cost and time. The site provides around-the-clock access to online tools, knowledge bases, and software. To access the Cisco TAC Web Site, go to this URL:

http://www.cisco.com/tac

All customers, partners, and resellers who have a valid Cisco service contract have complete access to the technical support resources on the Cisco TAC Web Site. The Cisco TAC Web Site requires a Cisco.com login ID and password. If you have a valid service contract but do not have a login ID or password, go to this URL to register:

http://www.cisco.com/register/

If you are a Cisco.com registered user, and you cannot resolve your technical issues by using the Cisco TAC Web Site, you can open a case online by using the TAC Case Open tool at this URL:

http://www.cisco.com/tac/caseopen

If you have Internet access, we recommend that you open P3 and P4 cases through the Cisco TAC Web Site.

Cisco TAC Escalation Center

The Cisco TAC Escalation Center addresses priority level 1 or priority level 2 issues. These classifications are assigned when severe network degradation significantly impacts business operations. When you contact the TAC Escalation Center with a P1 or P2 problem, a Cisco TAC engineer automatically opens a case.

To obtain a directory of toll-free Cisco TAC telephone numbers for your country, go to this URL:

http://www.cisco.com/warp/public/687/Directory/DirTAC.shtml

Before calling, please check with your network operations center to determine the level of Cisco support services to which your company is entitled: for example, SMARTnet, SMARTnet Onsite, or Network Supported Accounts (NSA). When you call the center, please have available your service agreement number and your product serial number.

Data Center Overview — Integrating Security, Load Balancing, and SSL Services using Service Modules

Data Centers, according to the report from the Renewable Energy Policy Project on Energy Smart Data Centers, are “an essential component of the infrastructure supporting the Internet and the digital commerce and electronic communication sector. Continued growth of these sectors requires a reliable infrastructure because … interruptions in digital services can have significant economic consequences”. According to the META Group, the average cost of an hour of downtime is estimated at $330,000. Strategic Research Corporation reports the financial impact of major outages is equivalent to US$6.5 million per hour for a brokerage operation, or US$2.6 million per hour for a credit-card sales authorization system.

Virtually every Enterprise has a Data Center, yet not every Data Center is designed to provide the proper levels of redundancy, scalability, and security. A Data Center design lacking in any of these areas is at some point going to fail to provide the expected services levels. Data Center downtime means the consumers of the information are not able to access it thus the Enterprise is not able to conduct business as usual.

Benefits of Building Data Centers

You can summarize the benefits of a Data Center in one sentence. Data Centers enable the consolidation of critical computing resources in controlled environments, under centralized management, that permit Enterprises to operate around the clock or according to their business needs. All Data Center services are expected to operate around the clock. When critical business applications are not available, the business is severely impacted and, depending on the outage, the company could cease to operate.

Building and operating Data Centers requires extensive planning. You should focus the planning efforts on those service areas you are supporting. High availability, scalability, security, and management strategies ought to be clear and explicitly defined to support the business requirements. Often times, however, the benefits of building Data Centers that satisfy such lists of requirements are better realized when the data center fails to operate as expected.

The loss of access to critical data is quantifiable and impacts the bottom line: revenue. There are a number of organizations that must address plans for business continuity by law, which include federal government agencies, financial institutions, healthcare and utilities. Because of the devastating effects of loss of data or access to data, all companies are compelled to look at reducing the risk and minimizing

the impact on the business. A significant portion of these plans is focused on Data Centers where critical business computing resources are kept. Understanding the impact of a Data Center failure in your Enterprise is essential. The following section introduces the Data Center role in the Enterprise network.

Data Centers in the Enterprise

Figure 1-1 presents the different building blocks used in the typical Enterprise network and illustrates the location of the Data Center within that architecture.

Figure 1-1 Enterprise Network Infrastructure

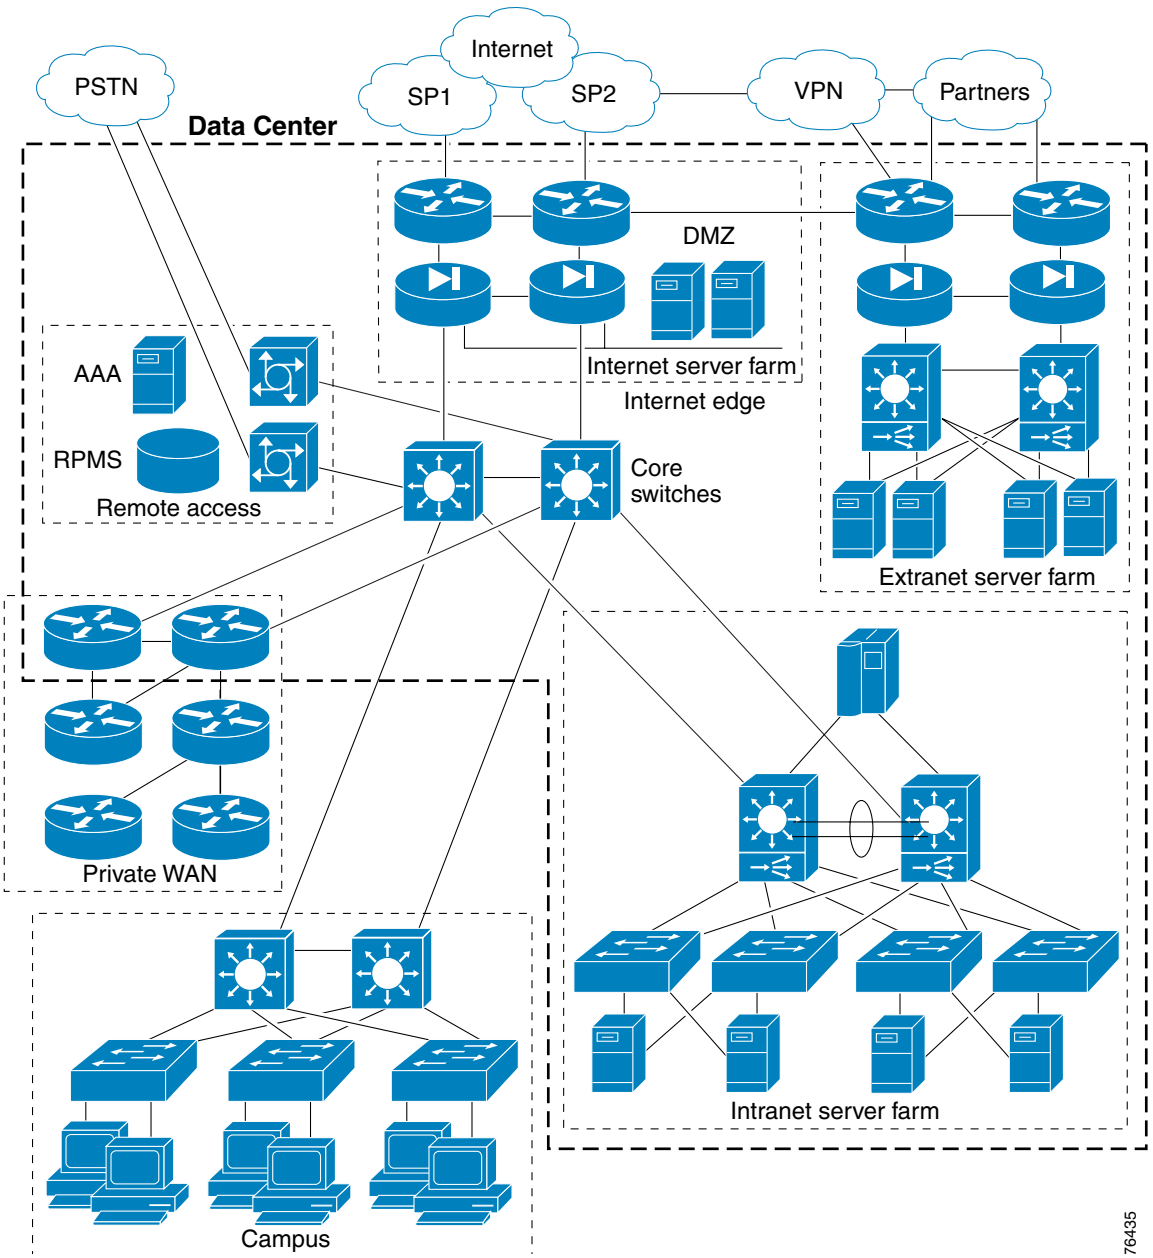

The building blocks of the typical Enterprise network include:

- Campus
- Private WAN
- Remote Access
- Internet server farm
- Extranet server farm
- Intranet server farm

Data Centers house many network infrastructure components that support the Enterprise network building blocks shown in Figure 1-1, such as the core switches of the Campus network or the edge routers of the Private WAN. Data Center designs however, include at least one type of server farm. These server farms may or may not be built as separate physical entities, depending on the business requirements of the Enterprise. For example, a single Data Center may use a shared infrastructure, resources such as servers, firewalls, routers, switches, etc., for multiple server farm types. Other Data Centers may require that the infrastructure for server farms be physically dedicated. Enterprises make these choices according to business drivers and their own particular needs. Once made, the best design practices presented in this chapter and subsequent design chapters can be used to design and deploy a highly available, scalable, and secured Data Center.

Data Center Architecture

The architecture of Enterprise Data Centers is determined by the business requirements, the application requirements, and the traffic load. These dictate the extent of the Data Center services offered, which translates into the actual design of the architecture. You must translate business requirements to specific goals that drive the detailed design. There are four key design criteria used in this translation process that help you produce design goals. These criteria are: availability, scalability, security, and management. Figure 1-2 shows the design criteria with respect to the Data Center architecture:

Figure 1-2 Architecture Layers and Design Criteria

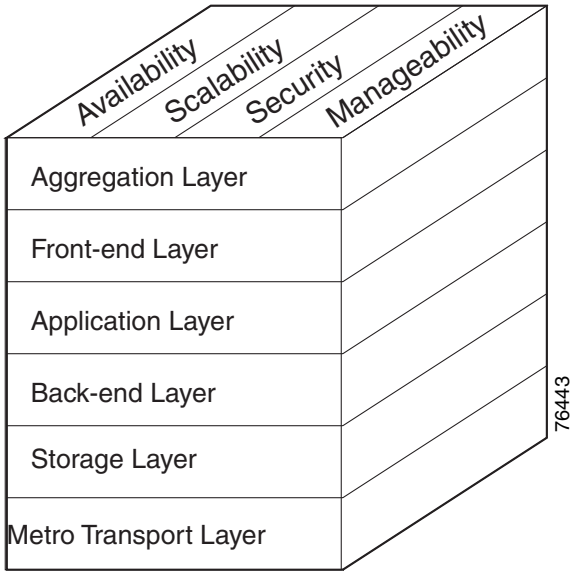

The purpose of using availability, scalability, security, and manageability as the design criteria is to determine what each layer of the architecture needs to meet the specific criteria. For instance, the answer to the question "how scalable the aggregation layer should be?" is driven by the business goals but is actually achieved by the Data Center design. Since the answer depends on which functions the aggregation layer performs, it is essential to understand what each layer does.

Your design goals and the services supported by the Data Center dictate the network infrastructure required. Figure 1-3 introduces the Data Center reference architecture.

Figure 1-3 Data Center Architecture

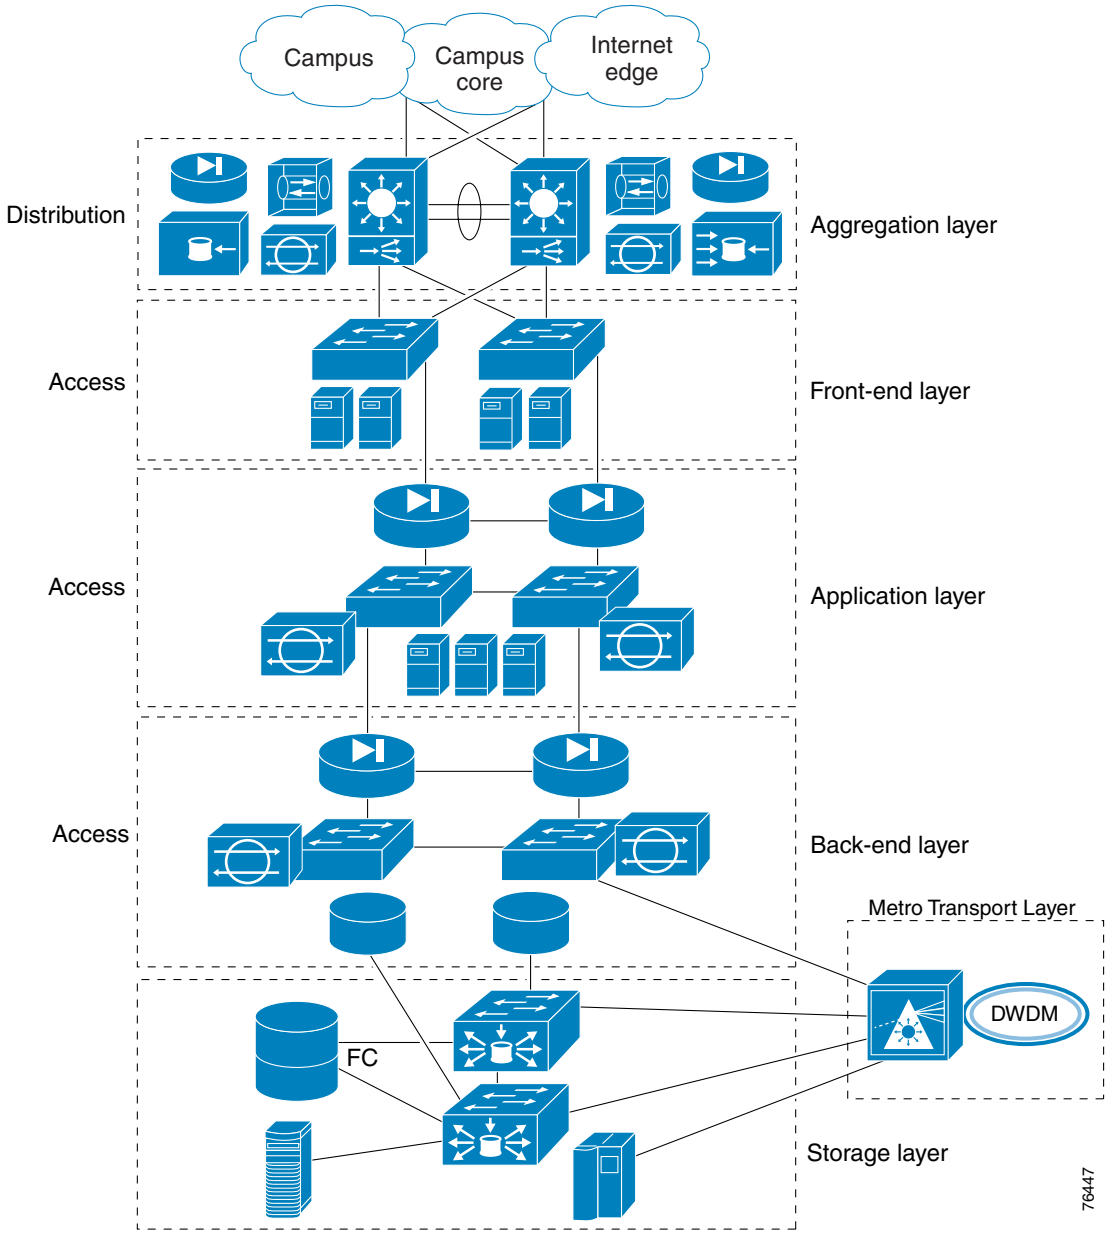

The architecture presents a layered approach to the Data Center design that supports N-Tier applications yet it includes other components related to other business trends. The layers of the architecture include:

- Aggregation
- Front-end
- Application
- Back-end
- Storage
- Metro Transport

Note The metro transport layer supports the metropolitan high-speed connectivity needs between distributed Data Centers.

The following sections provide a detailed description of these layers.

Aggregation Layer

The aggregation layer provides network connectivity between the server farms and the rest of the Enterprise network, provides network connectivity for Data Center service devices, and supports fundamental Layer 2 and Layer 3 functions. The aggregation layer is analogous to the campus network distribution layer. Data Center services that are common to servers in the front-end or other layers should be centrally located in the aggregation layer for predictability, consistency, and manageability. In addition to the multilayer switches (aggregation switches) that provide the Layer 2 and Layer 3 functionality, the aggregation layer includes, content switches, firewalls, IDSs, content engines, and SSL offloaders, as depicted in Figure 1-4.

Figure 1-4 Aggregation Layer

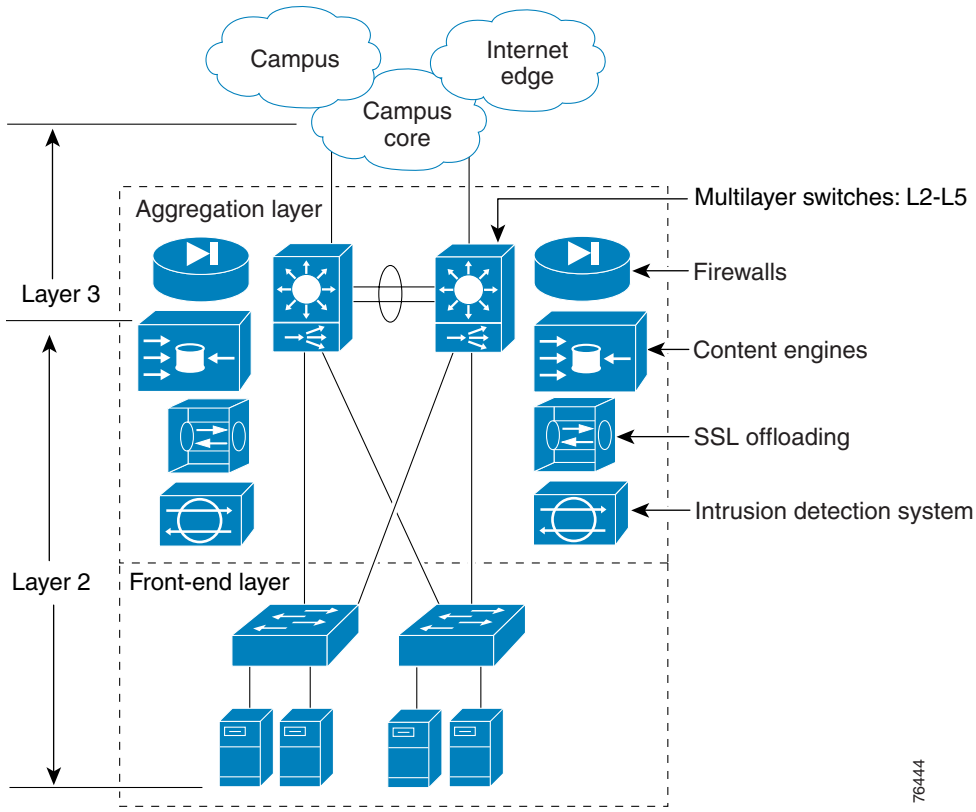

Front-End Layer

The front-end layer, analogous to the Campus access layer in its functionality, provides connectivity to the first tier of servers of the server farms. The front-end server farms typically include FTP, Telnet, TN3270, SMTP, Web servers, and other business application servers, in addition to network-based application servers, such as IPTV Broadcast servers, Content Distribution Managers, and Call Managers.

Specific features, such as Multicast and QoS that may be required, depend on the servers and their functions. For example, if live video streaming over IP is supported, multicast must be enabled; or if voice over IP is supported, QoS must be enabled. Layer 2 connectivity through VLANs is required between servers supporting the same application services for redundancy (dual homed servers on different Layer 2 switches), and between server and service devices such as content switches. Other requirements may call for the use of IDSs or Host IDSs to detect intruders or PVLANs to segregate servers in the same subnet from each other.

Application Layer

The application layer provides connectivity to the servers supporting the business logic, which are all grouped under the application servers tag. Applications servers run a portion of the software used by business applications and provide the communication logic between front-end and the back-end, which is typically referred to as the middleware or business logic. Application servers translate user requests to commands the back-end database systems understand.

The features required at this layer are almost identical to those needed in the front-end layer. Yet, additional security is typically used to tighten security between servers that face users and the next layer of servers, which implies firewalls in between. Additional IDSs may also be deployed to monitor different kinds of traffic types. Additional services may require load balancing between the web and application servers typically based on Layer 5 information, or SSL if the server-to-server communication is done over SSL. Figure 1-5 introduces the front-end, application, and back-end layers in a logical topology.

Figure 1-5 Front-End, Application, and Back-End Layers

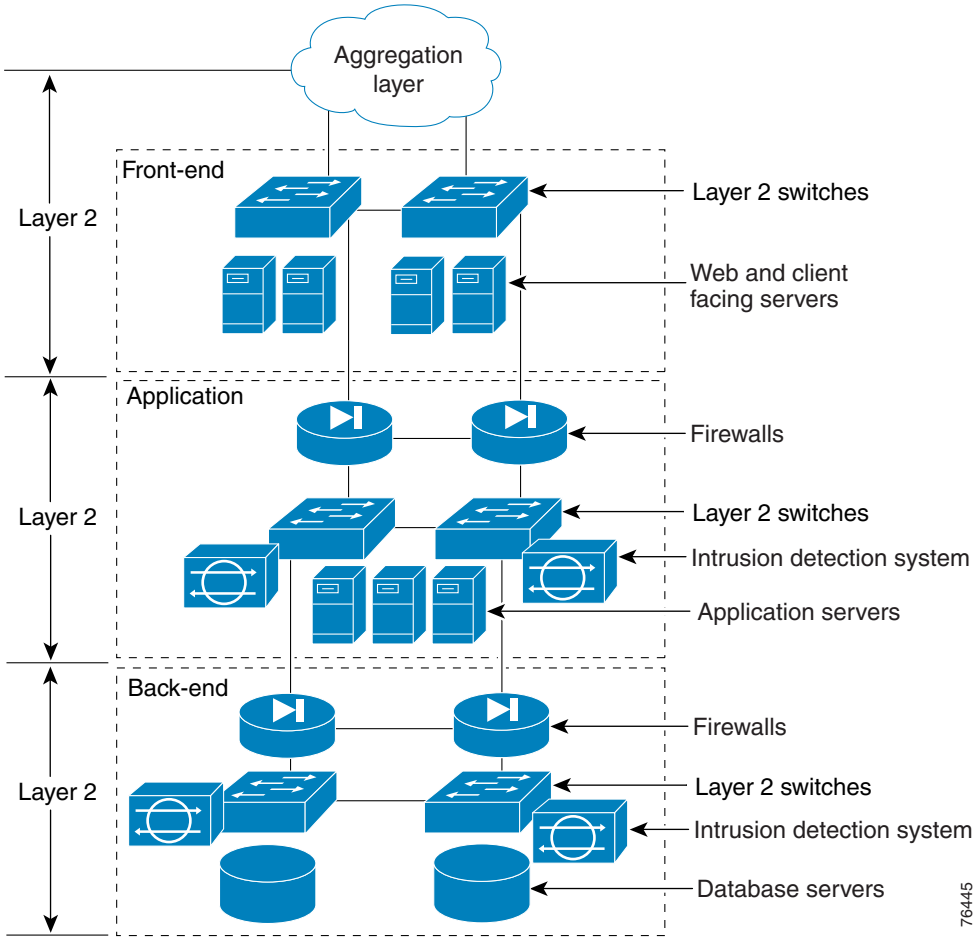

Back-End Layer

The back-end layer provides connectivity to the database servers. The feature requirements of this layer are almost identical to those of the application layer, yet the security considerations are more stringent and aimed at protecting the Enterprise data. The back-end layer is primarily for the relational database systems that provide the mechanisms to access the enterprise's information, which makes them highly critical. The hardware supporting the relational database systems range from medium sized servers to mainframes, some with locally attached disks and others with separate storage.

Storage Layer

The storage layer connects devices in the storage network using Fibre-Channel (FC) or iSCSI. The connectivity provided through FC switches is used for storage-to-storage communications between devices such as FC attached server and disk subsystems of tape units. iSCSI provides SCSI connectivity to servers over an IP network and is supported by iSCSI routers, port adaptors, and IP services modules. FC is typically used for block level access, whereas iSCSI is used for file level access.

Metro Transport Layer

The metro transport layer is used to provide a high speed connection between distributed Data Centers. These distributed Data Centers use metro optical technology to provide transparent transport media, which is typically used for database or storage mirroring and replication. This metro transport technology is also used for high speed campus-to-campus connectivity.

The high speed connectivity needs are either for synchronous or asynchronous communications, which depends on the recovery time expected when the primary data location fails. Disaster recovery and business continuance plans are the most common business driver behind the need for distributed Data Centers and the connectivity between them. Figure 1-6 presents a closer look to the logical view of the layer between the back-end and the metro transport.

Figure 1-6 Metro Transport Topology

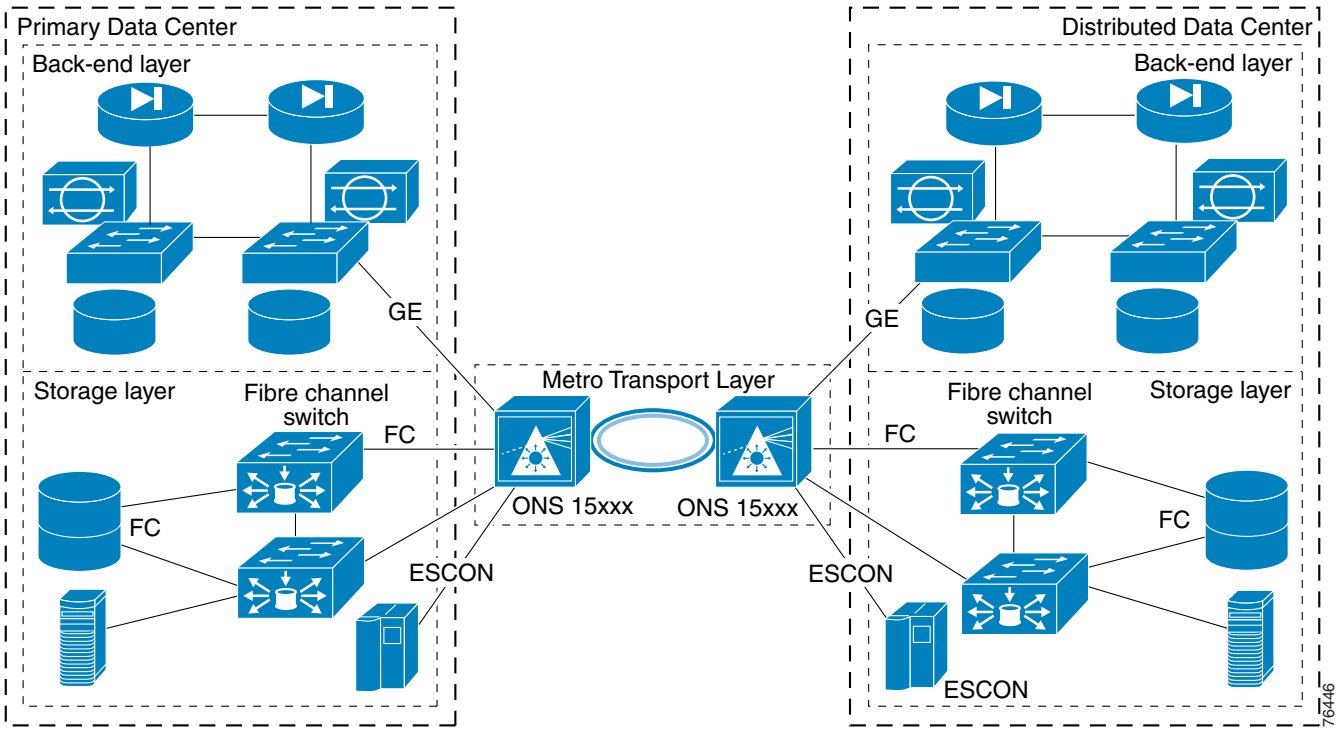

Distributed Data Centers

Distributed Data Centers provide redundancy for business applications. The primary Enterprise Data Center is a single point of failure when dealing with disasters. This could lead to application downtime leading to loss in productivity and lost business. Addressing this potentially high impact risk requires that the data is replicated at a remote location that acts as a backup or recovery site, the distributed Data Center, when the primary site is no longer operating.

The distributed Data Center, typically a smaller replica of the primary Data Center, takes over the primary data center responsibilities after a failure. With distributed Data Centers, data is replicated to the distributed Data Center over the metro transport layer. The clients are directed to the distributed Data Center when the primary Data Center is down. Distributed data centers reduce application down time for mission critical applications and minimize data loss.

Data Center Services

The Data Center is likely to support a number of services, which are the result of the application environment requirements. These services include:

- Infrastructure: Layer 2, Layer 3, Intelligent Network Services and Data Center Transport
- Application optimization services: content switching, caching, SSL offloading, And content transformation
- Storage: consolidation of local disks, Network Attached Storage, Storage Area Networks
- Security: access control lists, firewalls, and intrusion detection systems
- Management: Management devices applied to the elements of the architecture

The following section introduces the services details and their associated components.

Infrastructure Services

Infrastructure services include all core features needed for the Data Center infrastructure to function and serve as the foundation for all other Data Center services. The infrastructure features are organized as follows:

- Metro
- Layer 2
- Layer 3
- Intelligent Network Services

Metro Services

Metro services include a number of physical media access, such as Fibre-Channel and iSCSI, and metro transport technologies such as Dense Wave Division Multiplexing (DWDM), Coarse Wave Division Multiplexing (CWDM), SONET and 10GE. Metro transport technologies enable campus-to-campus and distributed Data Centers connectivity for a number of applications that require high bandwidth and low predictable delay. For instance, DWDM technology provides physical connectivity for a number of different physical media such as Gigabit Ethernet, ATM, Fibre Channel, and ESCON concurrently. Some instances where this connectivity is required are for long-haul Storage Area Networks (SAN) extension over SONET or IP and short-haul SAN extension over DWDM/CWDM, SONET, or IP (Ethernet).

Layer 2 Services

Layer 2 services support the Layer 2 adjacency between the server farms and the service devices, enable media access, provide transport technologies, and support a fast convergence, loop free, predictable, and scalable Layer 2 domain. In addition to LAN media access, such as Gigabit Ethernet, and ATM; there is

support for Packet over SONET (PoS), and IP over Optical media. Layer 2 domain features ensure the Spanning Tree Protocol (STP) convergence time for deterministic topologies is in the single digit seconds and that the failover and fallback scenarios are predictable. The list of features includes:

- 802.1s + 802.1w (Multiple Spanning-Tree)
- PVST+802.1w (Rapid Per VLAN Spanning-Tree)
- 802.3ad (Link Aggregate Control Protocol)
- 802.1q (trunking)
- LoopGuard
- Uni-Directional Link Detection (UDLD)
- Broadcast Suppression

Layer 3 Services

Layer 3 services enable fast convergence and a resilient routed network, including redundancy, for basic Layer 3 services, such as default gateway support. The purpose is to maintain a highly available Layer 3 environment in the Data Center where the network operation is predictable under normal and failure conditions. The list of available features includes:

- Static routing
- Border Gateway Protocol (BGP)
- Interior Gateway Protocols (IGPs): OSPF and EIGRP
- HSRP, MHSRP & VRRP

Intelligent Network Services

Intelligent network services include a number of features that enable applications services network wide. The most common features are QoS and Multicast. Yet there are other important intelligent network services, such as Private VLANs (PVLANs) and Policy Based Routing (PBR). These features enable applications, such as live or on demand video streaming and IP telephony, in addition to the classic set of enterprise applications. QoS in the Data Center is important for two reasons: marking, at the source, application traffic and port based rate limiting capabilities that enforces a proper QoS service class as traffic leaves the server farms. Multicast in the Data Center enables the capabilities needed to reach multiple users concurrently or servers to receive information concurrently (cluster protocols).

For more information on infrastructure services in the data center, see the *Data Center Networking: Infrastructure Architecture SRND*.

Application Optimization Services

Application optimization services include a number of features that provide intelligence to the server farms. These features permit the scaling of applications supported by the server farms and packet inspection beyond Layer 3 (Layer 4 or Layer 5).

The application services are:

- Server load balancing or content switching
- Caching
- SSL offloading

Content switching is used to scale application services by front ending servers and load balancing the incoming requests to those available servers. The load balancing mechanisms could be based on Layer 4 or Layer 5 information, thus allowing you to partition the server farms by the content they serve. For instance, a group of servers supporting video streaming could be partitioned on those that support MPEG versus the ones that support Quicktime or Windows Media. The content switch is able to determine the type of request, by inspecting the URL, and forwards it to the proper server. This process simplifies the management of the video servers and allows you to deal with scalability at a more granular level, per type of video server.

Caching, and in particular Reverse Proxy Caching, offloads the serving of static content from the server farms thus offloading CPU cycles, which increases scalability. The process of offloading occurs transparently for both the user and the server farm.

SSL offloading also offloads CPU capacity from the server farm by processing all the SSL traffic. The two key advantages to this approach are the centralized management of SSL services on a single device (as opposed to a SSL NIC per server) and the capability of content switches to load balance otherwise encrypted traffic in clear text.

For more information about application optimization services, see the *Data Center Networking: Optimizing Server and Application Environments SRND*.

Storage Services

Storage services include the storage network connectivity required for user-to-server and storage-to-storage transactions. The major features could be classified in the following categories:

- Network Attached Storage (NAS)
- Storage Area Networks (SAN) to IP: Fibre Channel and SCSI over IP
- Localized SAN fabric connectivity (Fibre Channel or iSCSI)
- Fibre Channel to iSCSI Fan-out

Storage consolidation leads to NAS and SAN environments. NAS relies on the IP infrastructure and, in particular, features such as QoS to ensure the proper file over the IP network to the NAS servers. SAN environments, commonly found in Data Centers, use Fibre Channel (FC) to connect servers to the storage device and to transmit SCSI commands between them. The SAN environments need to be accessible to the NAS and the larger IP Network.

FC over IP (FCIP) and SCSI over IP (iSCSI) are the emerging IETF standards that enable SCSI access and connectivity over IP. The transport of SCSI commands over IP enables storage-to-IP and storage-to-storage over an IP infrastructure.

SAN environments remain prevalent in Data Center environment, thus the localized SAN fabric becomes important to permit storage-to-storage block access communication at Fibre Channel speeds. There are other features focused on enabling FC to iSCSI fan-out for both storage-to-IP and storage-to-storage interconnects.

Security Services

Security services include a number of tools used in the application environment to increase security. The approach to security services in server farm environments is the result of increasing external threats but also internal attacks. This creates the need to have a tight security perimeter around the server farms and a plan to keep the security policies applied in a manner consistent with the risk and impact if the

Enterprise data was compromised. Since different portions of the Enterprise's data is kept at different tiers in the architecture, it is important to consider deploying security between tiers so that the specific tier has its own protection mechanisms according to likely risks.

Utilizing a layered security architecture provides a scalable modular approach to deploying security for the multiple data center tiers. The layered architecture makes use of the various security services and features to enhance security. The goal of deploying each of these security features and services is to mitigate against threats, such as:

- Unauthorized access
- Network reconnaissance
- IP spoofing
- Denial of Service
- Viruses and worms
- Layer 2 attacks

The security services offered in the data center include: access control lists (ACLs), firewalls, intrusion detection systems (IDS, Host IDS), authentication, authorization and accounting (AAA) mechanisms, and a number of other services that increase security in the data center.

ACLs

ACLs prevent unwanted access to infrastructure devices and, to a lesser extent, protect server farm services. You can apply ACLs at various points in the Data Center infrastructure. ACLs come in different types: Router ACLs (RACLs), VLAN ACLs (VACLs), and QoS ACLs. Each type of ACL is useful for specific purposes that, as their names indicate, are related to routers, VLANs, or QoS control mechanisms. An important feature of ACLs is the ability to perform packet inspection and classification without causing performance bottlenecks. This lookup process is possible when done in hardware, in which case the ACLs operate at the speed of the media, or at wire speed.

Firewalls

The placement of firewalls marks a clear delineation between highly secured and loosely secured network perimeters. While the typical location for firewalls remains the Internet edge and the edge of the Data Center, they are also used in multi-tier server farm environments to increase security between the different tiers.

Intrusion Detection

IDSs proactively address security issues. Intruder detection and the subsequent notification are a fundamental step to highly secure Data Centers where the goal is to protect the data. Host IDSs enable real-time analysis and reaction to hacking attempts on applications or Web servers. The Host IDS is able to identify the attack and prevent access to server resources before any unauthorized transactions occur.

AAA

AAA provides yet one more layer of security by preventing user access unless authorized, and by ensuring controlled user access to the network and network devices by a predefined profile. The transactions of all authorized and authenticated users are logged for accounting purposes, for billing, or for postmortem analysis.

Other Security Services

Additional security considerations may include the use of the following features or templates:

- One Time Passwords (OTPs)
- CDP to discover neighboring Cisco devices
- Default security templates for data center devices, such as routers, switches, firewalls and content switches
- SSH or IPSEC from user-to-device
- VTY security

For more information on security services, see the *Data Center Networking: Securing Server Farms SRND*.

Management Services

Management services refer to the ability to manage the network infrastructure that provides the support of all other services in the Data Center. The management of services in the Data Center include service provisioning, which depending on the specific service, requires its own set of management considerations. Each service is also likely supported by different organizational entities or even by distinct functional groups whose expertise is in the provisioning, monitoring, and troubleshooting of such service.

Cisco recommends that you have a network management policy in place that follows a consistent and comprehensive approach to managing Data Center services. Cisco follows the FCAPS OSI management standard and uses its management categories to provide management functionality. FCAPS is a model commonly used in defining network management functions and their role in a managed network infrastructure. The management features focus on the following categories:

- Fault management
- Configuration management
- Accounting management
- Performance management
- Security management

For more information on management services, see the *Data Center Networking: Optimizing Server and Application Environments SRND*.

Summary

The business requirements drive the application requirements, which in turn drive Data Center design requirements. The design process must take into account the current trends in application environments, such as the N-Tier model, to determine application requirements. Once application requirements are clear, the Data Center architecture needs to be qualified to ensure that its objectives are met and that application requirements are met.

A recommendation to the Data Center design process is that you consider the layers of the architecture that you need to support, given your specific applications, as the cornerstone of the services that you need to provide. These services must meet your objectives and must follow a simple set of design criteria to achieve those objectives. The design criteria include high availability, scalability, security, and management, which all together focus the design on the Data Center services.

Achieving your design goals translates to satisfying your application requirements and ultimately attaining your business objectives. Ensure that the Data Center design lets you achieve your current objectives, particularly as they relate to your mission critical applications. Knowing you can, enables you to minimize the business impact, as you would have quantified how resilient your Enterprise is to the always dynamic business conditions.

Integrating the Firewall Service Module

This chapter presents various deployment scenarios for the Firewall Services Module (FWSM) in the data center. The FWSM is a service module for the Catalyst 6500. The FWSM is a 5 Gigabit firewall based on the PIX code. The FWSM supports VLAN interfaces (100) and dynamic routing (OSPF).

Terminology

For the purpose of this chapter, a security domain is a collection of systems under a common security policy. A security domain can be made of multiple subnets and/or several server farms, where the server farm is a group of servers represented by a common Virtual IP address (VIP).

In this chapter, a Layer 3 VLAN means a VLAN that is not trunked to the access switches and is mainly used for communication between routing devices. A Layer 3 VLAN is carried on a single trunk in the network topology, specifically the trunk + channel that runs between the two aggregation switches.

A switched VLAN interface (SVI) is a VLAN interface defined on the MSFC. A VLAN configured on the Catalyst becomes an SVI when you use the interface vlan <vlan number> command to assign it an IP address. The creation of a VLAN by itself by the command "(config) vlan <vlan number>" does not create an SVI.

In the drawings that follow, the white box that contains the FWSM, the MSFC, and the load balancer represents a Catalyst 6500, and each component is basically a blade or a daughter card in the switch.

Overview

Data centers can take advantage of the FWSM to achieve the following goals:

- Control access to the intranet data center
- Create a demilitarized zone (DMZ) to host the Internet data center

In either scenario, you can decide how many security domains you want to create. You can use multiple security domains to either create multi-tier server farms or to just create multiple DMZs.

These main design categories can be further categorized based on the placement of the other network elements:

- The Multilayer Switching Feature Card (MSFC)
- Load balancer/s (Content Switching Module (CSM), Content Services Switches (CSS))

Note You are not required to use the MSFC in your design, nor you have to use a load balancer. When and if you decide to use the MSFC and/or a load balancing device in your data center, you will find that your design falls in one of the categories presented in this chapter.

The designs presented in this chapter take advantage of the MSFC for the routing. As a result the designs can be classified as:

- MSFC-outside
- MSFC-inside

Deployment Scenarios

The simplest design consists of using the FWSM to provide one single security domain in the intranet data center. This design is represented in Figure 2-1. On the left side of the picture, you see the physical diagram and on the right side, you see the logical diagram. The FWSMs are represented as external devices even if they are service modules inside the Catalyst 6500. Only two VLAN interfaces of the firewall are used: one for the inside and another one for the outside. In this design, the default gateway for the servers can be either the FWSM or a load balancing device, if present.

Figure 2-1 The FWSM in the Intranet Data Center

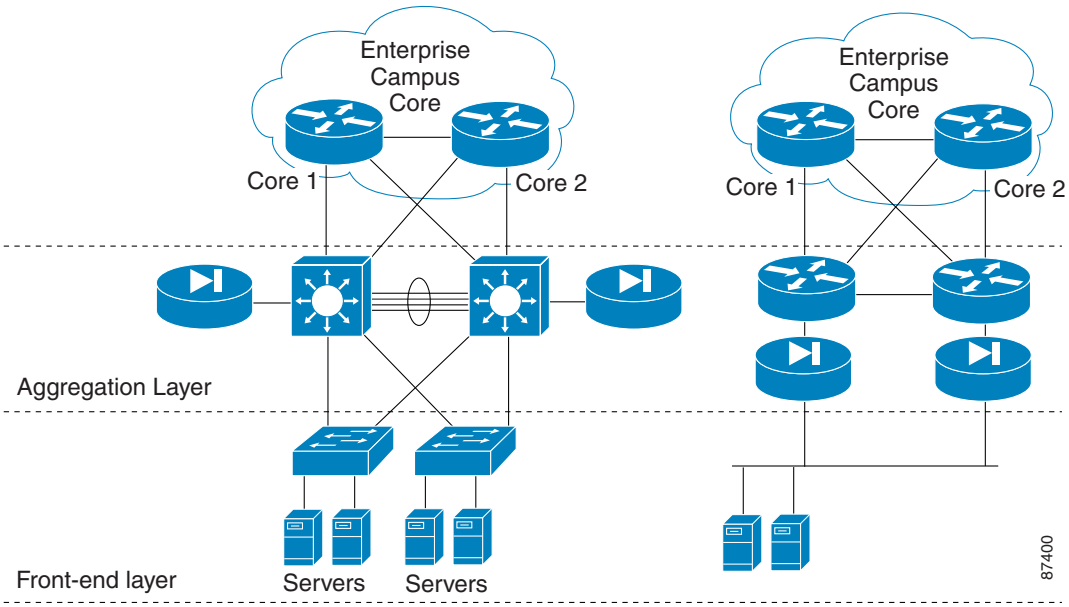

The second type of design (represented in Figure 2-2) is used to create a DMZ in the perimeter network. This is where you typically host your Internet data center.

On the left of the picture you can see the physical diagram and on the right you can see the logical diagram. When deploying the FWSM in the Internet edge, the typical connection to the Internet Service Provider (ISP) is through a pair of border routers. These border routers can be the same Catalyst 6500s hosting the FWSM or a separate pair of routers. In this design guide the Catalyst 6500s with FWSM are

not used as border routers, they just provide the aggregation layer for the Internet data center. You can decide how and if you want to use the MSFC. This design guide uses the MSFC to perform routing with the core of the enterprise. The default gateway for the servers in the DMZ is the FWSM.

Note If you attach the Catalyst 6500 switches with FWSM directly to the ISP network and make them the autonomous system border routers (ASBR) you have different options on how and if to use the MSFC. If you use the FlexWAN modules or the OSM modules, you have to place the MSFC facing the ISP and the FWSM on the inside because with these modules the traffic hits the MSFC first. If the ISP provides you with Gigabit attachment you have the choice of placing the MSFC on the outside or inside of the FWSM.

Figure 2-2 FWSM in the Internet Data Center

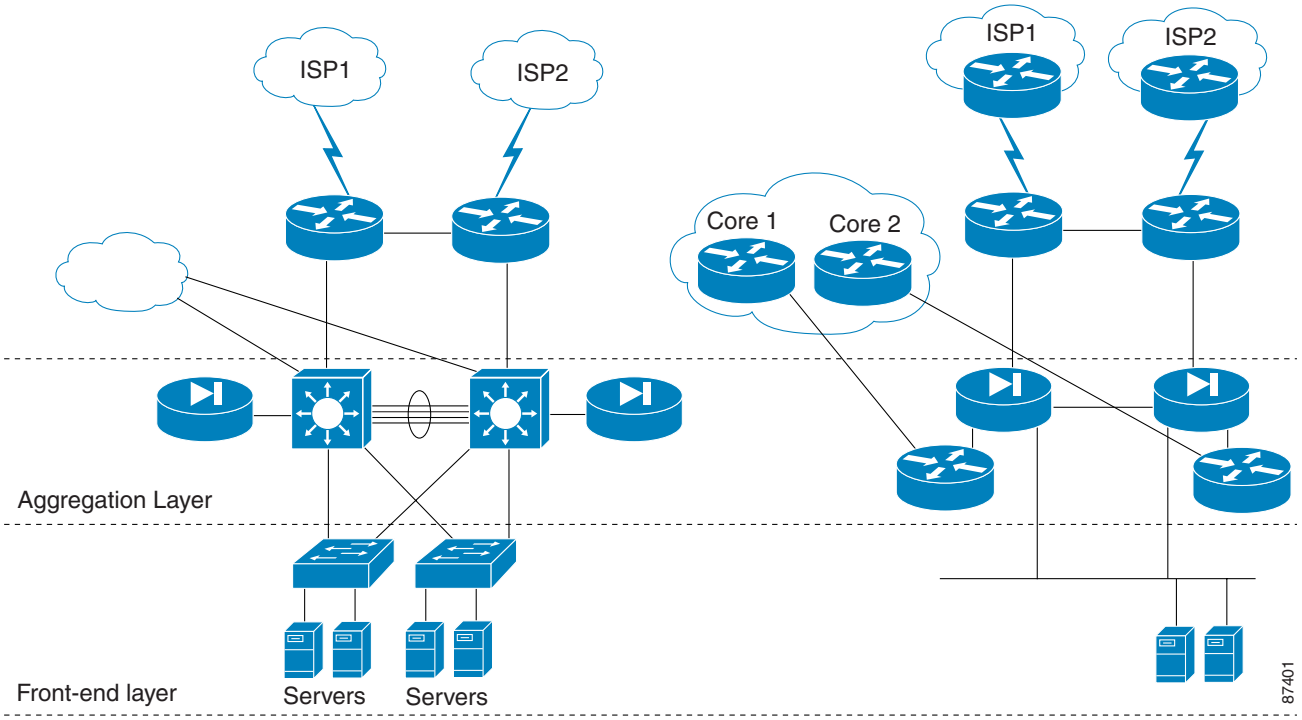

The FWSM can be used to segregate servers with different security levels. This is useful for servers that belong to different organizations or for applications to which you want to apply different filtering policies. When you want to segregate servers with different security levels, you must assign them to different VLANs. The FWSM uses VLANs as interfaces and you can assign a different security level to each of the VLANs. In Figure 2-3, the servers are assigned to two different segments. Each of these segments has an interface on the FWSM. The default gateway for the servers is the FWSM interface.

Figure 2-3 FWSM Used to Create Multiple Security Domains

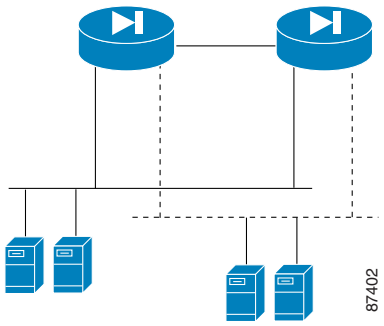

FWSM - MSFC Placement

One of the key elements that decide how the design works is the placement of the MSFC. The traffic hitting the aggregation switches from the core can hit the MSFC first and the FWSM second (MSFC-outside) or it can hit the FWSM first and the MSFC second (MSFC-inside). Typically the MSFC-outside design applies to the Intranet Data Center while the MSFC-inside applies to the Internet data center.

Note When deploying the FWSM you are not forced to place the MSFC somewhere in the network: the FWSM already provides you with OSPF routing, static routing and NAT functions. The use of the MSFC is dictated by needs such as terminating a BGP session, the use of FlexWAN or OSM cards, the need to run dynamic routing protocols such as EIGRP or IS-IS and more in general by routing requirements that cannot be accomplished with the FWSM. This design guide covers only designs that use the MSFC.

MSFC-Outside

The MSFC-outside design typically applies to an intranet data center. Placing the MSFC outside in the intranet data center means that the MSFC faces the core. There are multiple reasons for doing this, such as:

- The fact that the MSFC has more routing features
- The code is optimized to handle routing computations
- The MFSC is capable of dealing with bigger routing tables

For example, if you make the MSFC the area border router (ABR) in OSPF, you can limit the size of the routing table on the FWSM. You can have most of the routing recalculation happen on the MSFC and just propagate a default route to the firewall.

Having the MSFC as the router facing the core allows you to perform equal cost path load balancing on both Layer 3 uplinks that connect to the core. Having Layer 3 links to the core provides faster detection of a neighbor failure than having a shared segment.

With the MSFC-outside design, the default gateway for the servers is either the FWSM or the load balancer (such as the CSM).

In the case of an Internet data center, the MSFC-outside option is dictated by other factors such as the use of FlexWAN or OSM cards to connect to the Internet.

MSFC-Inside

The MSFC-inside design typically applies to the Internet data center. Placing the MSFC on the inside of the FWSM makes it possible for the MSFC to perform routing towards the enterprise core network. The FWSM provides routing to the border routers and the DMZ.

Using the FWSM facing the border routers requires having a shared segment between the aggregation switches: the two border routers both have an interface on this shared segment. If you want to load balance traffic to the border routers, you have to use Multigroup Hot Standby Router Protocol (MHSRP) on the interfaces of the routers facing the shared segment.

FWSM - CSM Placement

When attempting to provide load balancing and firewalling in the data center, you can choose whether you want to place the CSM outside the FWSM or on the inside of the FWSM. Both options are valid. When using the CSM on the inside, you can take advantage of the bridge mode to segregate VLANs of different security level consistently with the FWSM configuration. The result is that traffic from the core hits first the MSFC (MSFC-outside), then the FWSM, then the CSM. Figure 2-4 helps understanding the use of FWSM and CSM.

On the left of the picture, you can see the CSM operating in bridge mode between the servers and the FWSM, which means that the CSM bridges the server VLANs with the client VLANs. The advantage of using the CSM in bridge mode is that the FWSM performs the routing functions between the server VLANs. Server-to-server traffic for separate segments (such as from 10.20.5.x to 10.20.6.x) flows all the way to the FWSM and back to the CSM from the 10.20.6.x VLAN interface of the FWSM. The traffic from the 10.20.5.x servers going to the 10.20.6.x servers goes all the way to the FWSM and back to the CSM. The FWSM performs the routing and, the CSM performs the load balancing. In this design, the default gateway for the servers is the FWSM.

If you consider the fact that the CSM does not do any load balancing between the 10.20.5.x subnet and the 10.20.6.x unless the request for the Virtual IP address comes in from a FWSM interface, means that the design is equivalent to having multiple separate load balancers, one for each security domains. Figure 2-4 on the right, shows an equivalent design to the one with the shared CSM: one separate physical load balancer for each segment (security domain).

Figure 2-4 FWSM Used With a Shared CSM: Physical Diagram (Left), Logical Equivalent (Right)

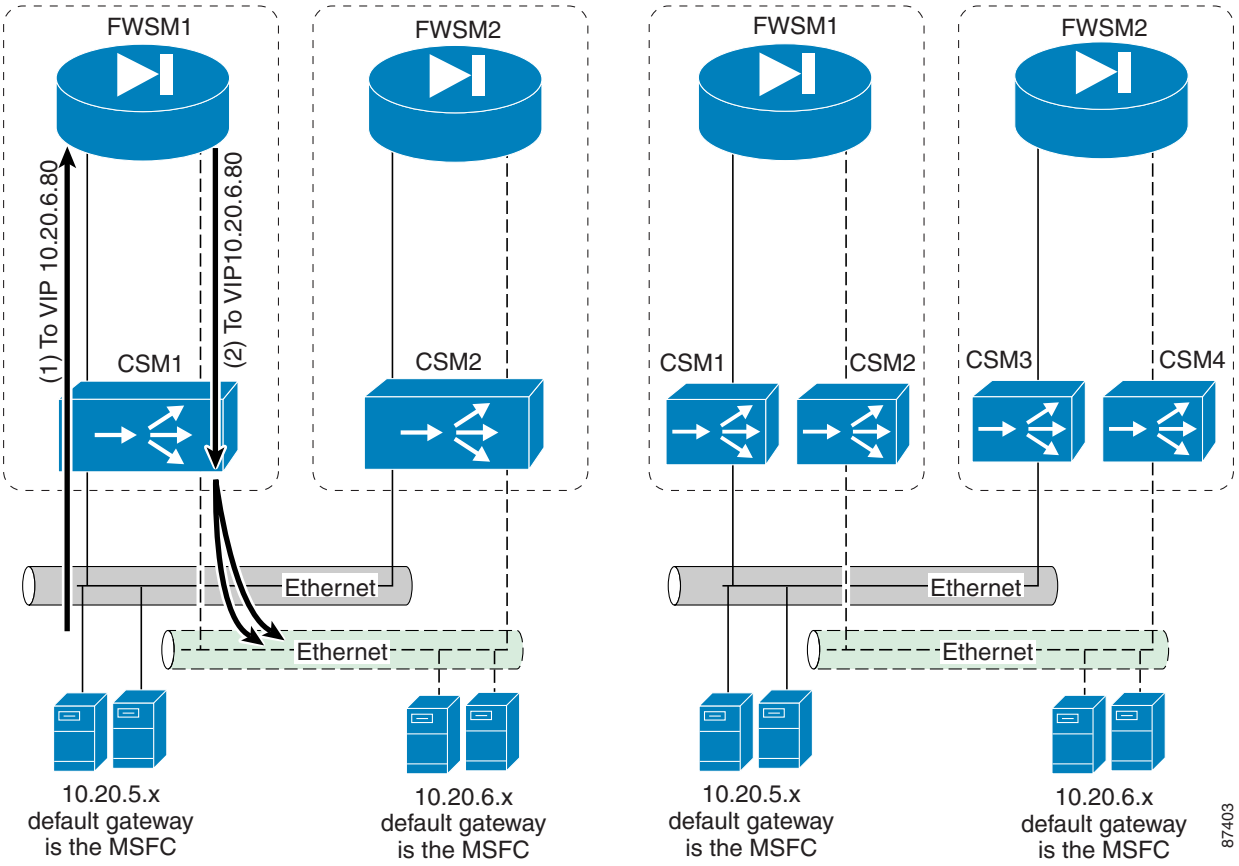

Redundancy

Deploying redundant FWSMs presents challenges very similar to deploying redundant CSMs. The FWSM operates in active/standby mode and provides stateful redundancy. The failover time is around 7s.

The communication between a redundant pair of FWSM uses a dedicated VLAN. This VLAN is trunked by the infrastructure switches. This approach requires at least some basic configuration on both the master and standby device in order for the election process to occur.

Both FWSMs in a redundant pair use the same MAC address when/if they are active. By doing this, there is no need to update the ARP tables of the adjacent routers when a failover happens.

On the FWSM, a command explicitly assigns the role for each device. Failover lan unit primary makes the firewall the primary device; similarly failover lan unit secondary makes the firewall the standby device.

The detection of a failure on the active unit is a combination of the following mechanisms:

- The active device sends a hello packet every 15s (this timer is configurable with the failover poll command and can be brought down to 3s). Hello packets are sent to all the interfaces.
- The standby unit monitors both the hello packets and the failover communication.

- Two consecutive missing hello packets trigger the failover tests.
- The failover tests consist in sending hello messages both on the interfaces and the failover connection. The units then monitor their interfaces to see if they have received traffic.
- There are additional tests the firewalls perform to decide which unit is faulty, which include an ARP test and a broadcast ping test.

The conclusion is that the convergence time by default is around 30s (twice the poll timer) and can be brought down to around 6s.

Configurations Description

Common Configurations: Layer 2/Layer 3

On the switch side, the only additional configuration that is required is the definition of which VLANs the switch needs to trunk to the FWSM. Use the firewall module and firewall vlan-group commands for this purpose. Notice that only one of the VLANs trunked to the FWSM is allowed to be an SVI.

Configuring VLANs

Perform the following steps on the switch side to configure the VLANs:

Step 1 Create the VLANs on the Catalyst 6000 (from the config-mode do “vlan <number>”), for example VLAN 20 and 30

Step 2 Trunk these VLANs between the aggregation Catalysts

Step 3 Define a VLAN-group for the FWSM: firewall vlan-group 1 20,30

Step 4 Assign the VLANs to a FWSM: firewall module <module number> vlan-group 1

Step 5 On the FWSM, assign names and security level to the VLAN interfaces. Use the nameif command.

- nameif vlan30 outside security0
- nameif vlan20 inside security100
- nameif <vlan #> <name> <security level>

Step 6 To monitor which VLANs are trunked between the Catalyst and the FWSM, use the show firewall module <module number> state command from the Catalyst console:

```
mp_agg2#sh firewall module 6 state
Firewall module 6:
 Switchport: Enabled
Administrative Mode: trunk
Operational Mode: trunk
Administrative Trunking Encapsulation: dot1q
Operational Trunking Encapsulation: dot1q
Negotiation of Trunking: Off
Access Mode VLAN: 1 (default)
Trunking Native Mode VLAN: 1 (default)
Trunking VLANs Enabled: 10,20,30,200
Pruning VLANs Enabled: 2-1001
Vlans allowed on trunk:10,20,30,200
Vlans allowed and active in management domain:10,20,30,200
Vlans in spanning tree forwarding state and not pruned:
```

```
   10,20,30,200
```

Step 7 To monitor which VLANs are configured, you can also issue the show vlan command from the FWSM CLI:

```
FWSM# sh vlan
10, 20, 30, 200
```

Configuring Trunks

When configuring the Catalyst 6500 with an integrated FWSM, remember to enable dot1q tagging for all the VLANs, including the native VLAN. You can do this by typing:

```
vlan dot1q tag native
```

Configuring IP Addresses

Only one of the VLANs listed under the firewall vlan-group command can be defined as a "vlan interface" (SVI) on the MSFC. For example, if the MSFC is on the outside, you can configure the following SVI:

```
interface Vlan30
 description FW-outide-vlan
 ip address 10.20.30.2 255.255.255.0
 ip ospf priority 10
!
```

On the firewall, assign IP addresses to both Vlan20 and Vlan30:

```
nameif vlan30 outside security0
nameif vlan20 inside security50
[…]
ip address outside 10.20.30.6 255.255.255.0
ip address inside 10.20.20.1 255.255.255.0
```

If you define in the vlan-group more than one SVI (Switched VLAN Interface) you see the following message:

```
mp_agg1(config)#firewall vlan-group 6 10,20
Found svi for vlan 10
Found svi for vlan 20
No more than one svi is allowed. Command rejected.
```

Use the no int vlan <vlan number> command to correct this problem. This command removes the SVI from the MSFC or changes the vlan-group list.

Configuring Routing

The FWSM can be configured to run OSPF. If the area is a totally stubby area, the configuration is as follows:

```
router ospf 20
        network 10.20.0.0 255.255.0.0 area 20
        area 20 stub no-summary
        log-adj-changes
!
```

Cisco recommends configuring the MSFC in such a way that the designated router (DR) is the SVI on the MSFC.

```
interface Vlan30
 description FW-outside-vlan
 ip address 10.20.30.2 255.255.255.0
 ip ospf priority 10
!
```

You can verify the routing by issuing the show route command:

```
FWSM# show route
eobc 127.0.0.0 255.255.255.0 127.0.0.61 1 CONNECT static
10.0.0.0 255.0.0.0 is variably subnetted, 9 subnets, 3 masks
C       10.20.30.0 255.255.255.0 is directly connected, outside
C       10.20.20.0 255.255.255.0 is directly connected, inside
O       10.21.0.12 255.255.255.252 0:42:54[110/11] via 10.20.30.3,
O       10.20.10.0 255.255.255.0 0:42:54[110/10] via 10.20.10.1,
O       10.21.0.8 255.255.255.252 0:42:54[110/11] via 10.20.30.3,
O       10.21.0.4 255.255.255.252 0:42:54[110/11] via 10.20.30.2,
O       10.20.3.0 255.255.255.0 0:42:54[110/11] via 10.20.30.2,
O       10.21.0.0 255.255.255.252 0:42:54[110/11] via 10.20.30.2,
O*IA 0.0.0.0 0.0.0.0 0:42:54[110/12] via 10.20.30.2, 0:42:54
                     [110/12] via 10.20.30.3
```

In some designs, you might need to configure redistribution of static routes on the FWSM. In this case, you need to configure the data center as an NSSA area. The following lines describe the configuration on the FWSM: the outside network is 10.20.30.x and the inside network is 10.20.5.x. The static route pushes traffic for 10.20.40.80 to the CSM on the inside interface of the FWSM.

```
router ospf 1
        network 10.20.5.0 255.255.255.0 area 20
        network 10.20.30.0 255.255.255.0 area 20
        area 20 nssa
        log-adj-changes
        redistribute static subnets
!
route inside 10.20.40.80 255.255.255.255 10.20.5.6 1
```

Configuring NAT

The following configuration allows an external client to have access to a server that is in the inside.

```
nameif vlan10 inside security100
nameif vlan171 outside security0
ip address inside 10.0.0.1 255.255.255.0
ip address outside 171.69.101.1 255.255.255.0
static(inside, outside) 171.69.101.4 10.0.0.4
```

The static command defines the higher security level interface (inside) to lower security level (outside) mapping and is followed by the public IP address and by the local IP address.

The following configuration allows internal clients to have access to the Internet.

```
nameif vlan10 inside security100
nameif vlan171 outside security0
ip address inside 10.0.0.1 255.255.255.0
ip address outside 171.69.101.1 255.255.255.0
global(outside) 2 171.69.101.5-171.69.101.14 netmask 255.255.255.0
nat(inside) 2 10.0.0.0 255.255.255.0
```

The **nat** command defines which IP addresses are eligible for NATing (local IP addresses). The global command defines the range of IP addresses to use as the pool. The number 2 used in the example binds the pool with the selected nat configuration.

Note In the Internet edge topology, it is common to define network address translation (NAT)at the edge of the infrastructure. It is also common and a recommended best practice to implement authentication between dynamic routing protocols at the edge of the network. In certain cases the authentication packets may be translated to another address which in turn may cause the authentication to fail. This is currently being researched and will be updated accordingly if configurations changes need to made.

Configuring Redundancy

The recommended configuration is with external redundancy: one FWSM per aggregation switch. One firewall is active, the other one is standby. You need to configure a separate VLAN for the failover protocol and trunk this VLAN between the two aggregation switches.

Steps on the Catalyst switches:

Step 1 Configure a VLAN on the Catalyst and use it only for the failover protocol, for example VLAN 200.

Step 2 Trunk this VLAN between the aggregation Catalysts.

Steps on the FWSM:

Step 1 Create a VLAN interface and give it a name, for example nameif vlan200 failover security99.

Step 2 Assign an IP address to VLAN 200 (called failover), for example ip address failover 10.20.200.1 255.255.255.0.

Step 3 Define VLAN 200 as the VLAN used by the failover protocol, for example failover lan interface failover.

Step 4 Define the firewall role (primary/ backup), for example failover lan unit primary.

Step 5 Define the IP addresses for the backup unit failover ip address.

Step 6 Define the link used for replication of the state information, for example failover link failover.

Step 7 Enable failover by typing failover.

The configuration is summarized below:

```
nameif vlan200 failover security99
ip address failover 10.20.200.1 255.255.255.0
failover lan unit primary
failover lan interface failover
failover timeout 0:00:00
failover poll 15
failover ip address outside 10.20.30.5
failover ip address inside 10.20.20.2
failover ip address failover 10.20.200.2
failover link failover
```

Intranet Data Center - One Security Domain

The single security domain configuration is characterized by having one single inside interface on the FWSM. Having the MSFC on the outside of the firewall lets the MSFC take care of the routing between the core and the data center.

Figure 2-5 FWSM with Single Security Domain and MSFC-Outside

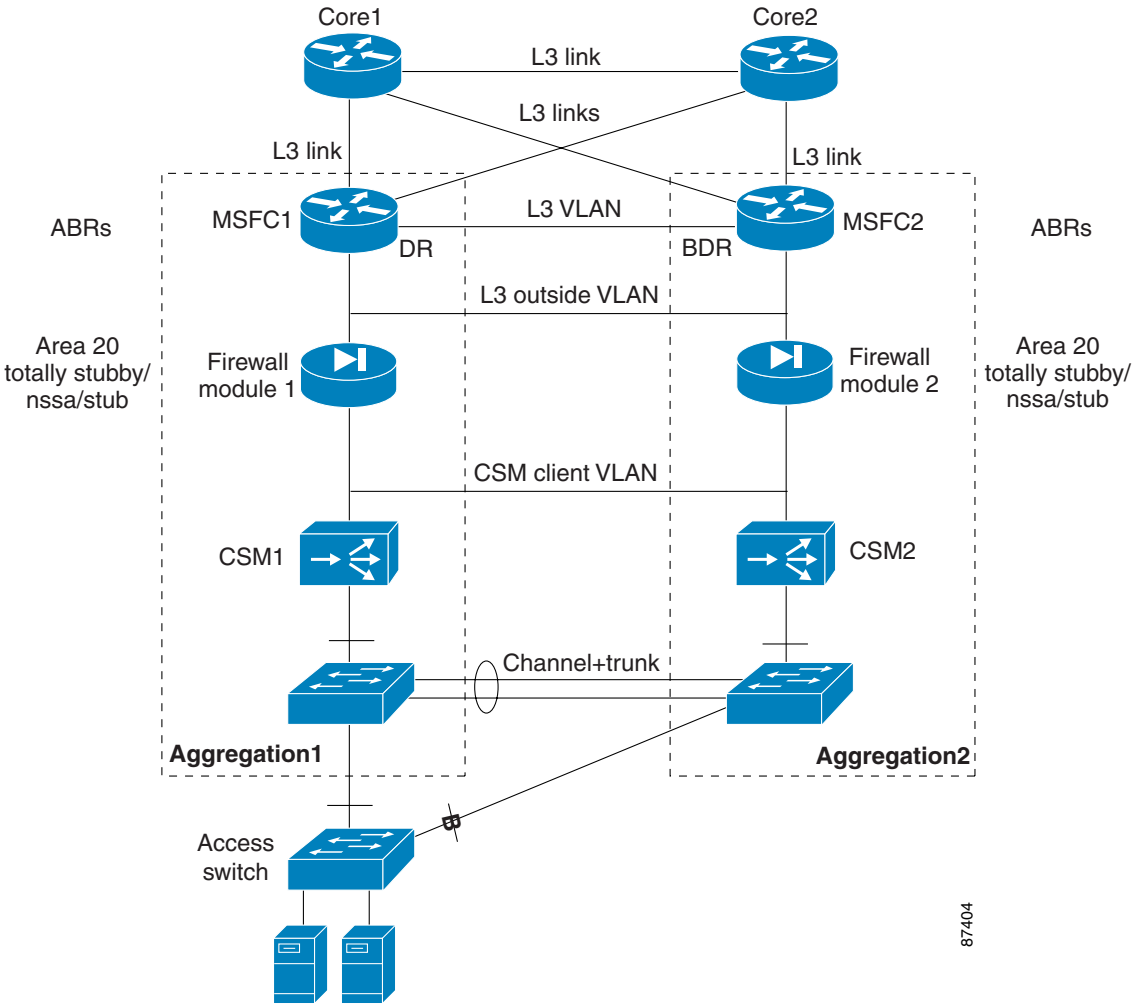

Because the MSFCs are outside, all the links to the core can be Layer 3 links. Equal paths achieve load balancing to the core routers. Also, the MSFC can be used as an ABR and advertises the summarized routes from the data center to the core. The area used for the data center can be a totally stubby, nssa, or stub area. The default gateway for the servers is either the load balancer or the firewall.

Internet Edge Deployment - MSFC-Inside

Figure 2-6 shows the deployment of the FWSM in the Internet edge. The MSFC-inside makes the MSFC available for routing to the core of the enterprise network. The default gateway for the servers is either the CSM or the FWSM. The FWSM shares a segment with the border routers. This common segment is bridged by the aggregation switches (outside VLAN in the picture) and provides connectivity between the FWSMs and the border routers.

In terms of routing, you can choose either static or dynamic routing. Dynamic routing has the advantage that you can dynamically advertise the default (or any other route) that you inject from the border routers. If you use OSPF, Cisco recommends making this area a not-so-stubby-area.

Figure 2-6 FWSM Design in the Internet Edge: MSFC Inside

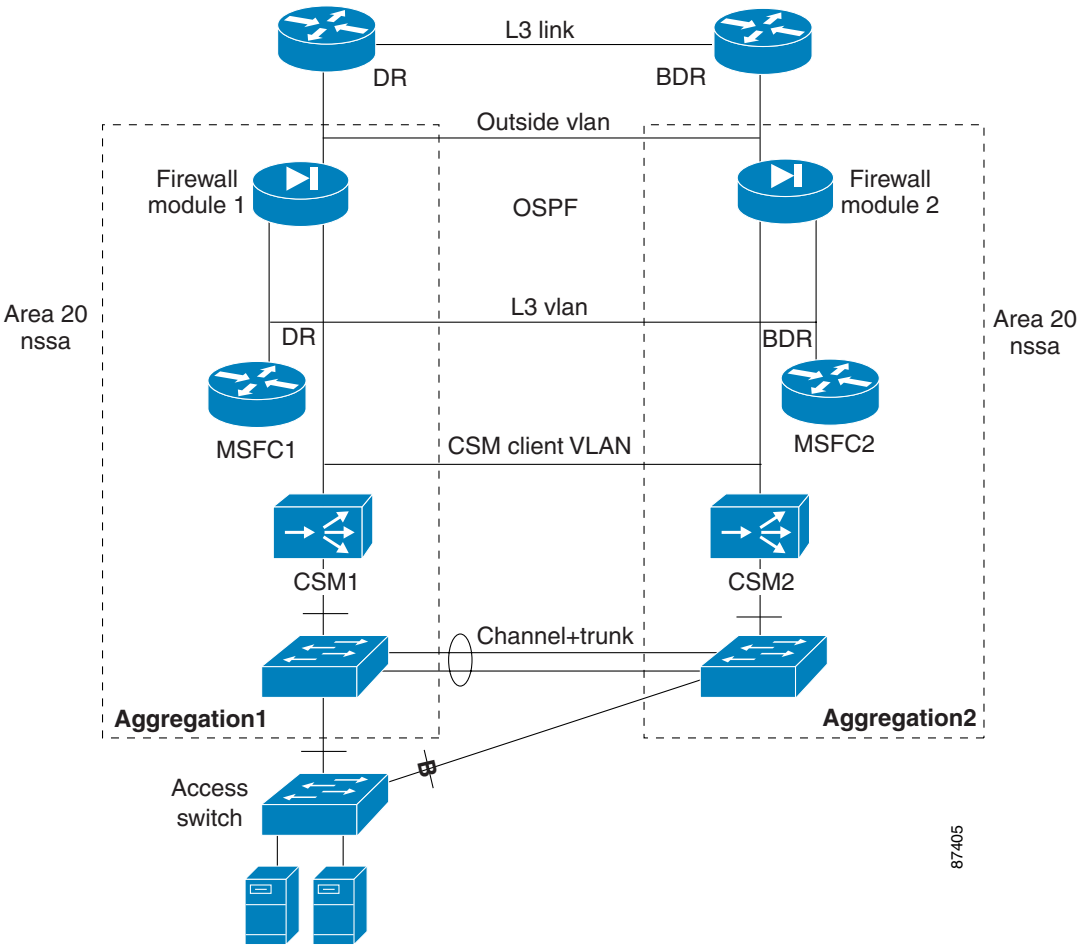

Multiple Security Domains / Multiple DMZs

A common requirement for data centers with multiple DMZs is to have the following traffic flow:

- From outside to DMZ1 (typically from clients to web servers)
- From DMZ1 to DMZ2 (typically from web servers to application servers or data base servers)

You do not typically want direct access from the outside network to DMZ2 with the above traffic pattern. As a result a possible configuration for the FWSM is the following one:

```
ip address outside 10.20.30.5 255.255.255.0
ip address dmz1 10.20.5.1 255.255.255.0
ip address dmz2 10.20.6.1 255.255.255.0
static (dmz1,outside) 10.20.5.0 10.20.5.0 netmask 255.255.255.0 0 0
static (dmz2,dmz1) 10.20.6.0 10.20.6.0 netmask 255.255.255.0 0 0
```

If you need to give direct access from the outside to DMZ2,you must configure an additional static NAT:

```
static (dmz2,outside) 10.20.6.0 10.20.6.0 netmask 255.255.255.0 0 0
```

For both scenarios, you need to configure ACLs. The configuration of ACLs is out of the scope of this chapter.

When configuring the data center for multiple security domains it is important to configure the CSM correctly. The following configuration achieves the behavior described in Figure 2-4. You need to configure the client and server side VLANs on the CSM and bridge them. The following is the configuration for Aggregation1, the configuration on Aggregation2 is similar with the exception of the highlighted fields:

```
module ContentSwitchingModule 5
vlan 5 client
  ip address 10.20.5.4 255.255.255.0
  alias 10.20.5.6 255.255.255.0
!
vlan 6 client
  ip address 10.20.6.4 255.255.255.0
  alias 10.20.6.6 255.255.255.0
!
vlan 10 server
ip address 10.20.5.4 255.255.255.0
!
vlan 12 server
  ip address 10.20.6.4 255.255.255.0
!
ft group 1 vlan 100
  priority 10
  heartbeat-time 5
 failover 4
!
```

Notice the following key points:

- In this example, the servers belong to two separate broadcast domains: 10.20.5.x and 10.20.6.x. You might not need to use two, you might just need one, in which case you would only bridge VLAN 5 with VLAN 10.
- Use the same IP address statement: “ip address 10.20.5.4" on both VLANs to bridge between VLAN5 and VLAN10.
- Use the same IP address statement: “ip address 10.20.6.4" to bridge between VLAN6 and VLAN12.

To complete the CSM configuration you need to configure vservers with the Virtual IP address and specify the incoming VLAN to match in the vserver. The reason for this is to enforce the FWSM as the entry point for each DMZ/security domain. For example, in Figure 2-4 the vserver for 10.20.6.80 needs to include the VLAN 6 as a matching criteria: VLAN 6 is shared between the CSM and FWSM.

The configuration looks like this:

```
vserver HTTP-VIP2
  virtual 10.20.6.80 tcp https
  vlan 6
  serverfarm WEB-VIP2
```

```
  persistent rebalance
  inservice
!
```

Configurations

These configurations show the deployment of the FWSM in an intranet data center, Internet data center and in an environment with multiple DMZs or security domains from the point of view of interoperability with the data center infrastructure.

Caution It is important to understand that the configurations in this chapter address the interoperability at Layer 2 and Layer 3, the access-list configurations should not be followed as implemented in this chapter because this is not a security document.

Intranet Data Center - One Security Domain

In this configuration, the Virtual IP address is 10.20.30.80. The FWSM provides translation between 10.20.30.80 and 10.20.5.80 (the VIP defined on the CSM). The MSFC advertises the 10.20.30.x subnet. The FWSM does not advertise the 10.20.5.x, but receives routing updates from the MSFC from the outside interface. If you want to advertise the 10.20.5.x subnet from the FWSM, you can modify the router OSPF configuration to include the network statement for this subnet.

Figure 2-7 Topology for the MSFC Outside Configuration

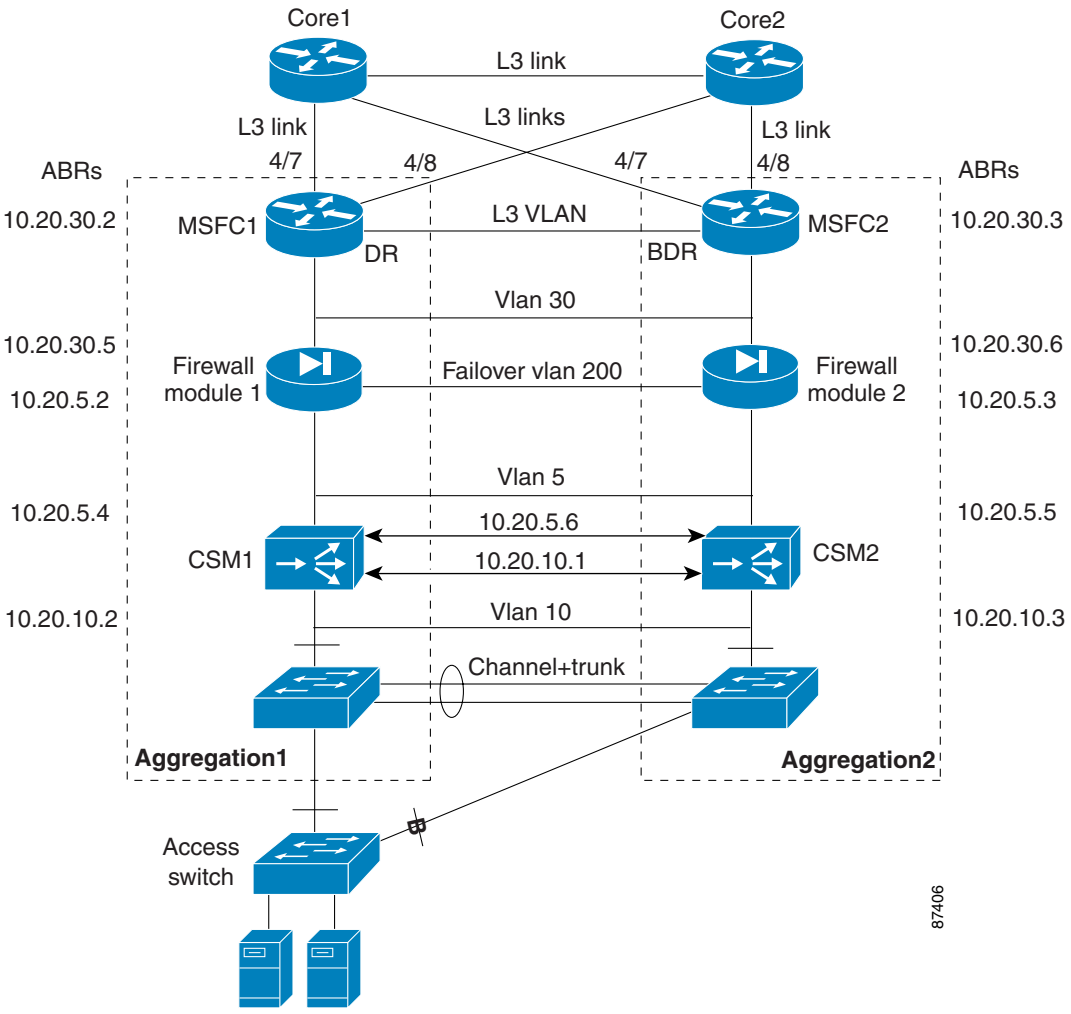

Aggregation1

```
!
version 12.1
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname mp_agg1
!
firewall module 6 vlan-group 6
firewall vlan-group 6  5,30,200
vtp domain mydomain
vtp mode transparent
ip subnet-zero
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
```

```
spanning-tree vlan 5,10,30,100,200 priority 8192
!
module ContentSwitchingModule 5
 vlan 5 client
  ip address 10.20.5.4 255.255.255.0
  alias 10.20.5.6 255.255.255.0
!
 vlan 10 server
  ip address 10.20.10.2 255.255.255.0
  alias 10.20.10.1 255.255.255.0
!
 probe TCP tcp
  interval 3
  failed 5
!
 serverfarm HTTP-SERVERS1
  nat server
  no nat client
  real 10.20.10.11
   inservice
  real 10.20.10.12
   inservice
  real 10.20.10.14
   inservice
  real 10.20.10.15
   inservice
!
 vserver HTTP-1
  virtual 10.20.5.80 tcp www
  serverfarm HTTP-SERVERS1
  persistent rebalance
  inservice
!
 ft group 1 vlan 100
  priority 20
  preempt
!
redundancy
 mode rpr-plus
 main-cpu
  auto-sync running-config
  auto-sync standard
!
vlan dot1q tag native
!
vlan 5
 name csm_client vlan
!
vlan 10
 name servers_group1
!
vlan 30
 name FW_outside
!
vlan 100
 name CSM_fault_tolerant
!
vlan 200
 name fw_failover_vlan
!
interface Loopback0
 ip address 3.3.3.3 255.255.255.255
!
interface Port-channel2
```

```
 no ip address
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 5,10,30,100,200,1002-1005
 switchport mode trunk
 spanning-tree guard loop
!
interface GigabitEthernet1/1
 description to_mp_agg2
 no ip address
 logging event link-status
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 5,10,30,100,200,1002-1005
 switchport mode trunk
 channel-group 2 mode active
 channel-protocol lacp
!
interface GigabitEthernet4/1
 description to_mp_acc1
 no ip address
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10, 1002-1005
 switchport mode trunk
!
interface GigabitEthernet4/6
 description to_mp_agg2
 no ip address
 logging event link-status
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 5,10,30,100,200,1002-1005
 switchport mode trunk
 channel-group 2 mode active
 channel-protocol lacp
!
interface GigabitEthernet4/7
 description to_mp_core1
 ip address 10.21.0.1 255.255.255.252
!
interface GigabitEthernet4/8
 description to_mp_core2
 ip address 10.21.0.5 255.255.255.252
!
interface Vlan30
 ip address 10.20.30.2 255.255.255.0
 no ip redirects
 ip ospf priority 10
!
router ospf 20
 log-adjacency-changes
 area 20 stub no-summary
 area 20 range 10.20.0.0 255.255.0.0
 network 10.20.0.0 0.0.255.255 area 20
 network 10.21.0.0 0.0.255.255 area 0
!
end
```

Aggregation2

```
!
version 12.1
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname mp_agg2
!
firewall module 6 vlan-group 6
firewall vlan-group 6  5,30,200
vtp domain mydomain
vtp mode transparent
ip subnet-zero
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
spanning-tree vlan 5,10,30,100,200 priority 16384
!
module ContentSwitchingModule 5
 vlan 5 client
  ip address 10.20.5.5 255.255.255.0
  alias 10.20.5.6 255.255.255.0
!
 vlan 10 server
  ip address 10.20.10.3 255.255.255.0
  alias 10.20.10.1 255.255.255.0
!
 probe TCP tcp
  interval 3
  failed 5
!
 serverfarm HTTP-SERVERS1
  nat server
  no nat client
  real 10.20.10.11
   inservice
  real 10.20.10.12
   inservice
  real 10.20.10.14
   inservice
  real 10.20.10.15
   inservice
!
 vserver HTTP-1
  virtual 10.20.5.80 tcp www
  serverfarm HTTP-SERVERS1
  persistent rebalance
  inservice
!
 ft group 1 vlan 100
  priority 10
  preempt
!
redundancy
 mode rpr-plus
 main-cpu
  auto-sync running-config
  auto-sync standard
!
vlan dot1q tag native
!
```

```
vlan 5
 name csm_vlan
!
vlan 10
 name servers-group1
!
vlan 30
 name FW_outside
!
vlan 100
 name CSM_fault_tolerant
!
vlan 200
 name fw_failover_vlan
!
interface Loopback0
 ip address 4.4.4.4 255.255.255.255
!
interface Port-channel2
 no ip address
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 5,10,30,100,200,1002-1005
 switchport mode trunk
!
interface GigabitEthernet1/1
 description to_mp_agg1
 no ip address
 logging event link-status
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 5,10,30,100,200,1002-1005
 switchport mode trunk
 channel-group 2 mode passive
 channel-protocol lacp
!
interface GigabitEthernet4/1
 description to_mp_acc1
 no ip address
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10, 1002-1005
 switchport mode trunk
!
interface GigabitEthernet4/6
 description to_mp_agg1
 no ip address
 logging event link-status
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 5,10,30,100,200,1002-1005
 switchport mode trunk
 channel-group 2 mode passive
 channel-protocol lacp
!
interface GigabitEthernet4/7
 description to_mp_core1
 ip address 10.21.0.9 255.255.255.252
!
interface GigabitEthernet4/8
 description to_mp_core2
 ip address 10.21.0.13 255.255.255.252
!
interface Vlan30
```

```
 description FW-outside-vlan
 ip address 10.20.30.3 255.255.255.0
 no ip redirects
 ip ospf priority 5
!
router ospf 20
 log-adjacency-changes
 area 20 stub no-summary
 area 20 range 10.20.0.0 255.255.0.0
 network 10.20.0.0 0.0.255.255 area 20
 network 10.21.0.0 0.0.255.255 area 0
!
```

FWSM1

```
FWSM1# wr t
Building configuration...
: Saved
:
FWSM Version 1.1(0)143
nameif vlan30 outside security0
nameif vlan5 inside security100
nameif vlan200 failover security50
enable password ZxdYhugKlKwPQ/MG encrypted
passwd ZxdYhugKlKwPQ/MG encrypted
hostname FWSM1
fixup protocol ftp 21
fixup protocol h323 H225 1720
fixup protocol h323 ras 1718-1719
fixup protocol ils 389
fixup protocol rsh 514
fixup protocol smtp 25
fixup protocol sqlnet 1521
fixup protocol sip 5060
fixup protocol skinny 2000
names
access-list incoming permit tcp any host 10.20.30.80 eq www
access-list incoming permit udp any any eq domain
access-list incoming permit icmp any host 10.20.30.80 echo
access-list incoming permit icmp any host 10.20.30.80 echo-reply
access-list outgoing permit tcp host 10.20.30.80 any eq www
access-list outgoing permit udp any any eq domain
access-list outgoing permit icmp host 10.20.30.80 any echo
access-list outgoing permit icmp host 10.20.30.80 any echo-reply
pager lines 24
icmp permit any outside
icmp permit any inside
mtu outside 1500
mtu inside 1500
mtu failover 1500
ip address outside 10.20.30.5 255.255.255.0
ip address inside 10.20.5.2 255.255.255.0
ip address failover 10.20.200.1 255.255.255.0
failover
failover lan unit primary
failover lan interface failover
failover timeout 0:00:00
failover poll 15
failover ip address outside 10.20.30.6
failover ip address inside 10.20.5.3
failover ip address failover 10.20.200.2
```

```
failover link failover
pdm history enable
arp timeout 14400
static (inside,outside) 10.20.30.80 10.20.5.80 netmask 255.255.255.255 0 0
access-group incoming in interface outside
access-group outgoing in interface inside
!
interface outside
!
interface inside
!
interface failover
!
router ospf 20
        network 10.20.0.0 255.255.0.0 area 20
        area 20 stub no-summary
        log-adj-changes
!
timeout xlate 3:00:00
timeout conn 1:00:00 half-closed 0:10:00 udp 0:02:00 rpc 0:10:00 h323 0:05:00 sip 0:30:
00 sip_media 0:02:00
timeout uauth 0:05:00 absolute
aaa-server TACACS+ protocol tacacs+
aaa-server RADIUS protocol radius
aaa-server LOCAL protocol local
no snmp-server location
no snmp-server contact
snmp-server community public
no snmp-server enable traps
floodguard enable
no sysopt route dnat
telnet timeout 5
ssh timeout 5
terminal width 80
Cryptochecksum:bb4e81430d01f89fb5e67a94445ed7c0
: end
[OK]
```

FWSM2

The FWSM2 configuration is the same with the exception of the following command:

```
failover lan unit secondary
```

Internet Edge Deployment - MSFC Inside

Figure 2-8 Topology for the Configuration in the Internet Edge

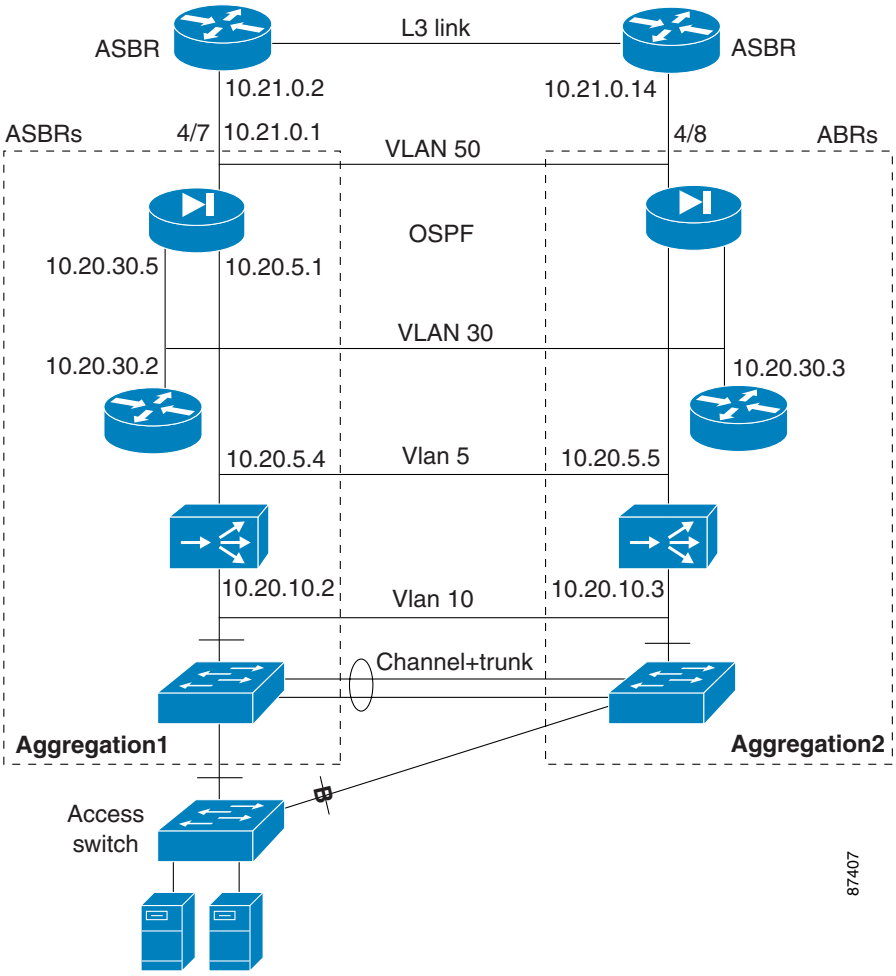

Aggregation1

```
!
version 12.1
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname mp_agg1
!
firewall module 6 vlan-group 6
firewall vlan-group 6  5,30,50,200
vtp domain mydomain
vtp mode transparent
ip subnet-zero
!
```

```
spanning-tree mode rapid-pvst
spanning-tree loopguard default
spanning-tree vlan 5,10,30,50,100,200 priority 8192
!
module ContentSwitchingModule 5
 vlan 10 server
  ip address 10.20.10.2 255.255.255.0
  alias 10.20.10.1 255.255.255.0
!
 vlan 5 client
  ip address 10.20.5.4 255.255.255.0
  alias 10.20.5.6 255.255.255.0
!
 probe TCP tcp
  interval 3
  failed 5
!
 serverfarm HTTP-SERVERS1
  nat server
  no nat client
  real 10.20.10.11
   inservice
  real 10.20.10.12
   inservice
  real 10.20.10.14
   inservice
  real 10.20.10.15
   inservice
  probe TCP
!
vserver HTTP-1
  virtual 10.20.5.80 tcp www
  serverfarm HTTP-SERVERS1
  persistent rebalance
  inservice
!
 ft group 1 vlan 100
  priority 20
  preempt
!
redundancy
 mode rpr-plus
 main-cpu
  auto-sync running-config
  auto-sync standard
!
vlan dot1q tag native
!
vlan 5
 name csm_vlan
!
vlan 10
 name servers_group1
!
vlan 30
 name FW inside
!
vlan 50
 name outside vlan
!
vlan 100
 name CSM_fault_tolerant
!
vlan 200
```

```
 name fw_failover_vlan
!
interface Loopback0
 ip address 3.3.3.3 255.255.255.255
!
interface Port-channel2
 no ip address
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 5,10,30,50,100,200,1002-1005
 switchport mode trunk
 spanning-tree guard loop
!
interface GigabitEthernet1/1
 description to_mp_agg2
 no ip address
 logging event link-status
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 5,10,30,50,100,200,1002-1005
 switchport mode trunk
 channel-group 2 mode active
 channel-protocol lacp
!
interface GigabitEthernet4/1
 description to_mp_acc1
 no ip address
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10,1002-1005
 switchport mode trunk
!
interface GigabitEthernet4/6
 description to_mp_agg2
 no ip address
 logging event link-status
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 5,10,30,50,100,200,1002-1005
 switchport mode trunk
 channel-group 2 mode active
 channel-protocol lacp
!
interface GigabitEthernet4/7
 description to_mp_core1
 no ip address
 switchport
 switchport access vlan 50
!
interface Vlan30
 ip address 10.20.30.2 255.255.255.0
 no ip redirects
 ip ospf priority 10
!
router ospf 20
 log-adjacency-changes
 area 20 nssa
 network 10.20.0.0 0.0.255.255 area 20
!
```

Aggregation2

```
!
version 12.1
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname mp_agg2
!
firewall module 6 vlan-group 6
firewall vlan-group 6  5,30,50,200
vtp domain mydomain
vtp mode transparent
ip subnet-zero
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
spanning-tree vlan 5,10,30,50,100,200 priority 16384
!
module ContentSwitchingModule 5
 vlan 5 client
  ip address 10.20.5.5 255.255.255.0
  alias 10.20.5.6 255.255.255.0
!
 vlan 10 server
  ip address 10.20.10.3 255.255.255.0
  alias 10.20.10.1 255.255.255.0
!
 probe TCP tcp
  interval 3
  failed 5
!
 serverfarm HTTP-SERVERS1
  nat server
  no nat client
  real 10.20.10.11
   inservice
  real 10.20.10.12
   inservice
  real 10.20.10.14
   inservice
  real 10.20.10.15
   inservice
!
 vserver HTTP-1
  virtual 10.20.5.80 tcp www
  serverfarm HTTP-SERVERS1
  persistent rebalance
  inservice
!
 ft group 1 vlan 100
  priority 10
  preempt
!
!
redundancy
 mode rpr-plus
 main-cpu
  auto-sync running-config
  auto-sync standard
!
vlan dot1q tag native
```

```
!
vlan 5
 name csm_vlan
!
vlan 10
 name servers-group1
!
vlan 30
 name FW_inside
!
vlan 50
 name outside vlan
!
vlan 100
 name CSM_fault_tolerant
!
vlan 200
 name ft_vlan
!
interface Loopback0
 ip address 4.4.4.4 255.255.255.255
!
interface Port-channel2
 no ip address
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 5,10,30,50,100,200,1002-1005
 switchport mode trunk
!
interface GigabitEthernet1/1
 description to_mp_agg1
 no ip address
 logging event link-status
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 5,10,30,50,100,200,1002-1005
 switchport mode trunk
 channel-group 2 mode passive
 channel-protocol lacp
!
interface GigabitEthernet4/1
 description to_mp_acc1
 no ip address
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10,1002-1005
 switchport mode trunk
!
interface GigabitEthernet4/6
 description to_mp_agg1
 no ip address
 logging event link-status
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 5,10,30,50,100,200,1002-1005
 switchport mode trunk
 channel-group 2 mode passive
 channel-protocol lacp
!
interface GigabitEthernet4/8
 description to_mp_core2
 no ip address
 switchport
 switchport access vlan 50
```

```
!
interface Vlan30
 description FW-inside-vlan
 ip address 10.20.30.3 255.255.255.0
 no ip redirects
 ip ospf priority 5
!
router ospf 20
 log-adjacency-changes
 area 20 nssa
 network 10.20.0.0 0.0.255.255 area 20
!
```

FWSM1

```
FWSM1# sh run
: Saved
:
FWSM Version 1.1(0)143
nameif vlan5 dmz security10
nameif vlan30 inside security100
nameif vlan50 outside security0
nameif vlan200 failover security50
enable password ZxdYhugKlKwPQ/MG encrypted
passwd ZxdYhugKlKwPQ/MG encrypted
hostname FWSM1
fixup protocol ftp 21
fixup protocol h323 H225 1720
fixup protocol h323 ras 1718-1719
fixup protocol ils 389
fixup protocol rsh 514
fixup protocol smtp 25
fixup protocol sqlnet 1521
fixup protocol sip 5060
fixup protocol skinny 2000
names
access-list dmzout permit ip any any
access-list incoming permit ip any any
access-list outgoing permit ip any any
pager lines 24
mtu inside 1500
mtu outside 1500
mtu failover 1500
ip address dmz 10.20.5.1 255.255.255.0
ip address inside 10.20.30.5 255.255.255.0
ip address outside 10.21.0.1 255.255.255.0
ip address failover 10.20.200.1 255.255.255.0
failover
failover lan unit primary
failover lan interface failover
failover timeout 0:00:00
failover poll 15
failover ip address dmz 10.20.5.2
failover ip address inside 10.20.30.6
failover ip address outside 10.21.0.13
failover ip address failover 10.20.200.2
pdm history enable
arp timeout 14400
global(outside) 1 10.21.0.20 10.21.0.40 netmask 255.255.0.0
nat(inside) 1 0.0.0.0 0.0.0.0
```

```
static (dmz,outside) 10.20.5.0 10.20.5.0 netmask 255.255.0.0 0 0
access-group outgoing in interface inside
access-group incoming in interface outside
access-group dmzout in interface dmz
!
interface inside
!
interface outside
!
interface failover
!
router ospf 20
        network 10.20.0.0 255.255.0.0 area 20
        network 10.21.0.0 255.255.0.0 area 20
        area 20 nssa
        log-adj-changes
```

Note For the ease of configuration in this deployment guide only, one OSPF process was utilized. The recommended ESE configuration requires two different OSPF processes at the edge of the network in this topology. This allows for you to control the outbound advertised routes accordingly in the architecture by defining route-maps to allow only specific segments or VIP addresses to be advertised. If only one OSPF process were to be used, all of the directly connected interfaces or routes in the FWSM would be propagated upstream. In certain environments, this could become a vulnerability. Refer to the *Data Center Networking: Internet Edge Design Architectures SRND* for more information.

```
!
timeout xlate 3:00:00
timeout conn 1:00:00 half-closed 0:10:00 udp 0:02:00 rpc 0:10:00 h323 0:05:00 sip 0:30:
00 sip_media 0:02:00
timeout uauth 0:05:00 absolute
aaa-server TACACS+ protocol tacacs+
aaa-server RADIUS protocol radius
aaa-server LOCAL protocol local
no snmp-server location
no snmp-server contact
snmp-server community public
no snmp-server enable traps
floodguard enable
no sysopt route dnat
telnet timeout 5
ssh timeout 5
terminal width 80
Cryptochecksum:f27eaa0b20da06fe1d7353fcc1e97683
: end
```

FWSM2

The FWSM2 configuration is the same with the exception of the following command:

```
failover lan unit secondary
```

Multiple Security Domains - Shared Load Balancer

Figure 2-9 Design with CSM in Bridge Mode

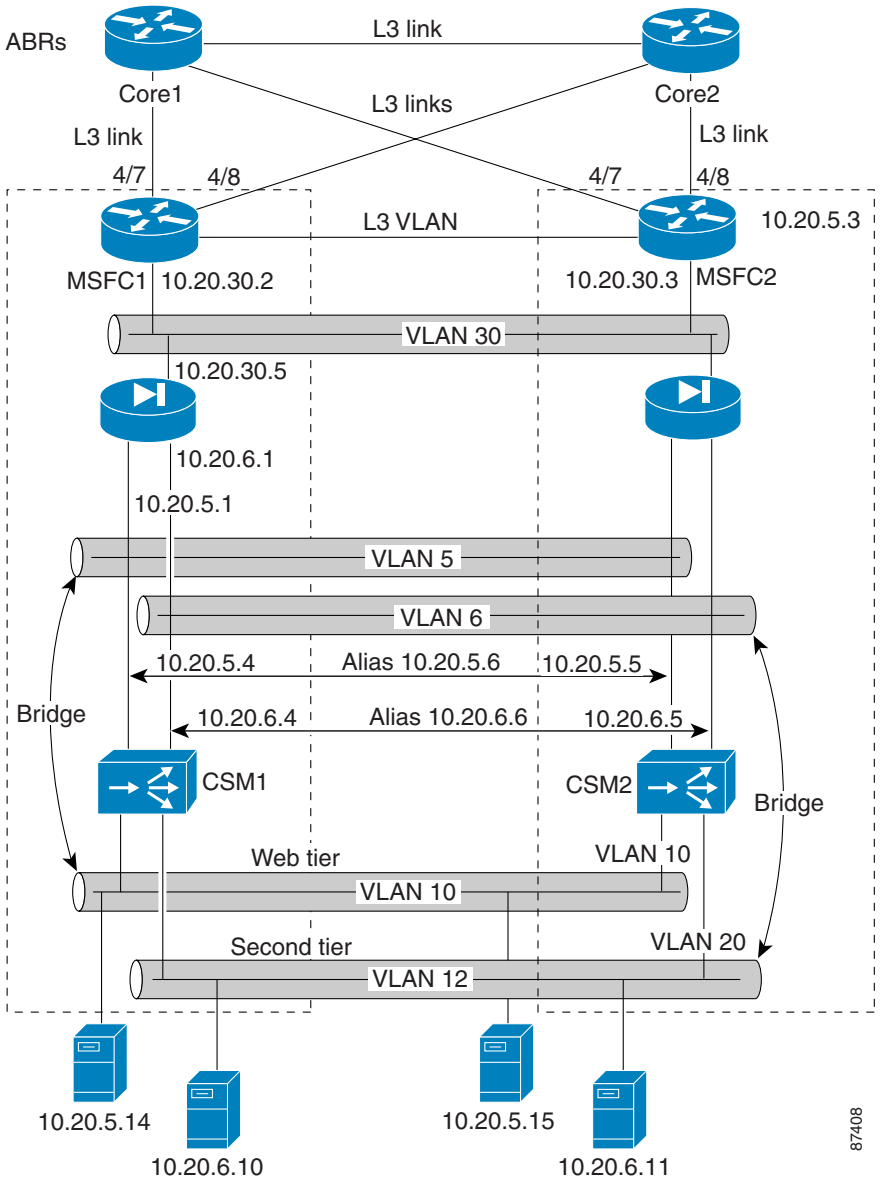

Aggregation1

```
!
version 12.1
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
```

```
hostname mp_agg1
!
firewall module 6 vlan-group 6
firewall vlan-group 6  5,6,30,200
vtp domain mydomain
vtp mode transparent
ip subnet-zero
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
spanning-tree vlan 5,6,10,12,30,100,200 priority 8192
!
module ContentSwitchingModule 5
 vlan 5 client
  ip address 10.20.5.4 255.255.255.0
  alias 10.20.5.6 255.255.255.0
!
vlan 6 client
  ip address 10.20.6.4 255.255.255.0
  alias 10.20.6.6 255.255.255.0
!
 vlan 10 server
  ip address 10.20.5.4 255.255.255.0
!
 vlan 12 server
  ip address 10.20.6.4 255.255.255.0
!
probe TCP tcp
  interval 2
  failed 3
!
 probe ICMP icmp
  interval 3
  failed 5
!
serverfarm WEB-VIP1
  nat server
  no nat client
  real 10.20.5.14
   inservice
  real 10.20.5.15
   inservice
  probe TCP
!
 serverfarm WEB-VIP2
  nat server
  no nat client
  real 10.20.6.10
   inservice
  real 10.20.6.11
   inservice
  probe TCP
!
vserver HTTP-VIP1
  virtual 10.20.5.80 tcp www
  vlan 5
  serverfarm WEB-VIP1
  persistent rebalance
  inservice
!
vserver HTTP-VIP2
  virtual 10.20.6.80 tcp www
  vlan 6
  serverfarm WEB-VIP2
```

```
  persistent rebalance
  inservice
!
ft group 1 vlan 100
  priority 20
  preempt
!
redundancy
 mode rpr-plus
 main-cpu
  auto-sync running-config
  auto-sync standard
!
vlan dot1q tag native
!
vlan 5
 name WEB1VLAN
!
vlan 6
 name WEB2VLAN
!
vlan 10
 name WEB1SERVERVLAN
!
vlan 12
 name WEB2SERVERVLAN
!
vlan 30
 name outside VLAN
!
vlan 100
 name CSM_fault_tolerant
!
vlan 200
 name FWSM fault tolerant
!

interface Port-channel2
 no ip address
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 5,6,10,12,30,100,200,1002-1005
 switchport mode trunk
 spanning-tree guard loop
!
interface GigabitEthernet1/1
 description to_mp_agg2
 no ip address
 logging event link-status
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 5,6,10,12,30,100,200,1002-1005
 switchport mode trunk
 channel-group 2 mode active
 channel-protocol lacp
!
interface GigabitEthernet4/1
 description to_mp_acc1
 no ip address
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10,12,1002-1005
 switchport mode trunk
!
```

```
interface GigabitEthernet4/6
 description to_mp_agg2
 no ip address
 logging event link-status
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 5,6,10,12,30,100,200,1002-1005
 switchport mode trunk
 channel-group 2 mode active
 channel-protocol lacp
!
interface GigabitEthernet4/7
 description to_mp_core1
 ip address 10.21.0.1 255.255.255.252
!
interface GigabitEthernet4/8
 description to_mp_core2
 ip address 10.21.0.5 255.255.255.252
!
interface Vlan30
 ip address 10.20.30.2 255.255.255.0
 no ip redirects
 ip ospf priority 10
 standby 1 ip 10.20.30.1
 standby 1 priority 110
 standby 1 preempt
!
router ospf 20
 log-adjacency-changes
 area 20 stub no-summary
 network 10.20.0.0 0.0.255.255 area 20
 network 10.21.0.0 0.0.255.255 area 20
!
```

Aggregation2

```
!
version 12.1
service timestamps debug uptime
service timestamps log uptime
no service password-encryption
!
hostname mp_agg2
!

firewall module 6 vlan-group 6
firewall vlan-group 6  5,6,30,200
vtp domain mydomain
vtp mode transparent
ip subnet-zero
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
spanning-tree vlan 5,6,10,12,30,100,200 priority 16384
!
module ContentSwitchingModule 5
vlan 5 client
  ip address 10.20.5.5 255.255.255.0
  alias 10.20.5.6 255.255.255.0
!
vlan 6 client
```

```
  ip address 10.20.6.5 255.255.255.0
  alias 10.20.6.6 255.255.255.0
!
vlan 10 server
  ip address 10.20.5.5 255.255.255.0
!
vlan 12 server
  ip address 10.20.6.5 255.255.255.0
!
probe TCP tcp
  interval 2
  failed 3
!
 probe ICMP icmp
  interval 3
  failed 5
!
serverfarm WEB-VIP1
  nat server
  no nat client
  real 10.20.5.14
   inservice
  real 10.20.5.15
   inservice
  probe TCP
!
 serverfarm WEB-VIP2
  nat server
  no nat client
  real 10.20.6.10
   inservice
  real 10.20.6.11
   inservice
  probe TCP
!
vserver HTTP-VIP1
  virtual 10.20.5.80 tcp www
  vlan 5
  serverfarm WEB-VIP1
  persistent rebalance
  inservice
!
vserver HTTP-VIP2
  virtual 10.20.6.80 tcp www
  vlan 6
  serverfarm WEB-VIP2
  persistent rebalance
  inservice
!
 ft group 1 vlan 100
  priority 10
  preempt
!
!
redundancy
 mode rpr-plus
 main-cpu
  auto-sync running-config
  auto-sync standard
!
vlan dot1q tag native
!
vlan 5
 name WEB1VLAN
```

```
!
vlan 6
 name WEB2VLAN
!
vlan 10
 name WEB1SERVERVLAN
!
vlan 12
 name WEB2SERVERVLAN
!
vlan 30
 name outside VLAN
!
vlan 100
 name CSM_fault_tolerant
!
vlan 200
 name FWSM fault tolerant
!
interface Port-channel2
 no ip address
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 5,6,10,12,30,100,200,1002-1005
 switchport mode trunk
!
interface GigabitEthernet1/1
 description to_mp_agg1
 no ip address
 logging event link-status
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 5,6,10,12,30,100,200,1002-1005
 switchport mode trunk
 channel-group 2 mode passive
 channel-protocol lacp
!
interface GigabitEthernet4/1
 description to_mp_acc1
 no ip address
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 10,12,1002-1005
 switchport mode trunk
!
interface GigabitEthernet4/6
 description to_mp_agg1
 no ip address
 logging event link-status
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 5,6,10,12,30,100,200,1002-1005
 switchport mode trunk
 channel-group 2 mode passive
 channel-protocol lacp
!
interface GigabitEthernet4/7
 description to_mp_core1
 ip address 10.21.0.9 255.255.255.252
!
interface GigabitEthernet4/8
 description to_mp_core2
 ip address 10.21.0.13 255.255.255.252
!
```

```
interface Vlan30
 ip address 10.20.30.3 255.255.255.0
 no ip redirects
 ip ospf priority 5
 standby 1 ip 10.20.30.1
 standby 1 priority 100
 standby 1 preempt
!
router ospf 20
 log-adjacency-changes
 area 20 stub no-summary
 network 10.20.0.0 0.0.255.255 area 20
 network 10.21.0.0 0.0.255.255 area 20
!
FWSM1
FWSM Version 1.1(0)143
nameif vlan30 outside security0
nameif vlan5 dmz1 security10
nameif vlan6 dmz2 security20
nameif vlan200 failover security50
enable password ZxdYhugKlKwPQ/MG encrypted
passwd ZxdYhugKlKwPQ/MG encrypted
hostname FWSM1
fixup protocol ftp 21
fixup protocol h323 H225 1720
fixup protocol h323 ras 1718-1719
fixup protocol ils 389
fixup protocol rsh 514
fixup protocol smtp 25
fixup protocol sqlnet 1521
fixup protocol sip 5060
fixup protocol skinny 2000
names
access-list dmz1out permit ip any any
access-list dmz2out permit ip any any
access-list incoming permit ip any any
pager lines 24
icmp permit any outside
icmp permit any inside
icmp permit any dmz1
icmp permit any dmz2
mtu outside 1500
mtu inside 1500
mtu failover 1500
ip address outside 10.20.30.5 255.255.255.0
ip address dmz1 10.20.5.1 255.255.255.0
ip address dmz2 10.20.6.1 255.255.255.0
ip address failover 10.20.200.1 255.255.255.0
failover
failover lan unit primary
failover lan interface failover
failover timeout 0:00:00
failover poll 15
failover ip address outside 10.20.30.6
failover ip address dmz1 10.20.5.2
failover ip address dmz2 10.20.6.2
failover ip address failover 10.20.200.2
pdm history enable
arp timeout 14400
static (dmz1,outside) 10.20.5.0 10.20.5.0 netmask 255.255.255.0 0 0
static (dmz2,dmz1) 10.20.6.0 10.20.6.0 netmask 255.255.255.0 0 0
```

Note If you need to give access from the outside to DMZ2, you need to add a static NAT configuration, static(dmz2, outside) 10.20.6.0 10.20.6.0 netmask 255.255.255.0 0 0

```
access-group incoming in interface outside
access-group dmz1out in interface dmz1
access-group dmz2out in interface dmz2
!
interface outside
!
interface inside
!
interface failover
!
router ospf 20
        network 10.20.0.0 255.255.0.0 area 20
        area 20 stub no-summary
        log-adj-changes
!
timeout xlate 3:00:00
timeout conn 1:00:00 half-closed 0:10:00 udp 0:02:00 rpc 0:10:00 h323 0:05:00 sip 0:30:
00 sip_media 0:02:00
timeout uauth 0:05:00 absolute
aaa-server TACACS+ protocol tacacs+
aaa-server RADIUS protocol radius
aaa-server LOCAL protocol local
no snmp-server location
no snmp-server contact
snmp-server community public
no snmp-server enable traps
floodguard enable
no sysopt route dnat
telnet timeout 5
ssh timeout 5
terminal width 80
Cryptochecksum:3fbc14aa231b4db698a4cf1f7ae1fc64
: end
[OK]
```

FWSM2

The FWSM2 configuration is the same with the exception of the following command:

```
failover lan unit secondary
```

Integrating the Content Switching Module

Content switching is fast becoming a widely accepted technology for scaling and optimizing application services residing in the data center. Content switching deployments utilize intelligent services to maximize the efficiency of the individual servers residing in the server farm, which, in turn, increases the overall performance and resiliency of the application services located on these servers. The Content Switching Module (CSM) introduces a high performance module-based solution for scaling Layer 4 and Layer 5 data center server farm services.

Overview

What is the CSM

The CSM is a content switch on a blade available for the Catalyst 6500 or 7600 platforms. It is capable of maintaining multiple Gigabit throughput rates (4 Gbps max inbound and outbound) while providing Layer 4-7 inspection and load balancing services.

The CSM provides performance for approximately 75,000-80,000 cps (connections per second) when configured to function at Layer 5, and approximately 125,000 cps when configured to function at Layer 4.

CSM Requirements

The CSM requires the use of Native IOS in the Catalyst switch and does not function with CatOS. Native IOS requires the presence of an MSFC on the Supervisor card. The following table presents the software and hardware recommendations for the CSM and the Catalyst switch.

Table 3-1 CSM hardware and software requirements

CSM Software Release	Hardware Part Number	Software Part Number	Hardware Requirements	IOS Release
3.1(1a)	WS-X6066-SLB-APC	SC6K-3.1-CSM	Supervisor 1A w/ MSFC & PFC -Or- Supervisor 2 Module with MSFC 2	12.1(13)E

Interoperability Details

This section discusses CSM interoperability in the data center architecture.

Data Center Network Infrastructure

In order to provide optimal content delivery through the use of expanded Layer 4-7 services on the CSM, the content must first be accessible through lower layer network protocols. As a prerequisite to adding content switching services to the data center, the existence of a highly available and scalable Layer 2 and Layer 3 network infrastructure is of utmost importance. Figure 3-1 illustrates a data center infrastructure designed around these requirements.

Figure 3-1 Layer 2 and Layer 3 Data Center Network Infrastructure

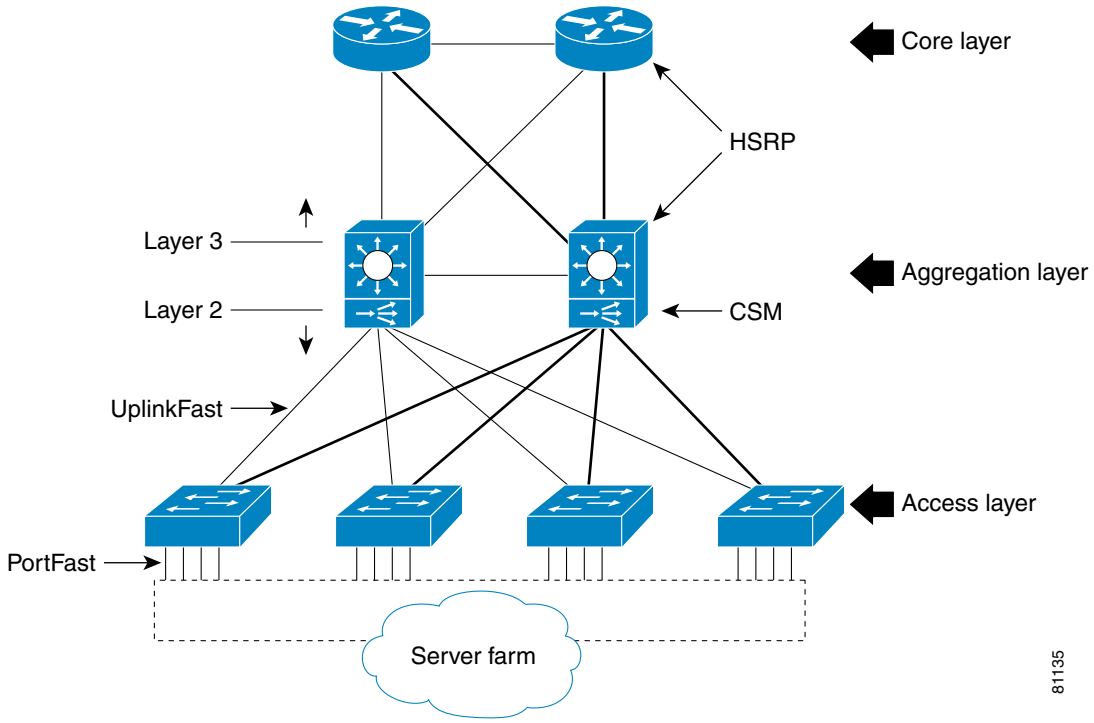

By leveraging existing Layer 2 and Layer 3 protocols on Cisco switches and routers, it is possible to provide a network infrastructure that supports high availability and low convergence times, which in turn, increases the likelihood that any service failures are made transparent to the end user. For detailed information regarding data center network infrastructure deployment, see the Data Center Networking: Infrastructure Design Architecture SRND located at:

http://www.cisco.com/en/US/netsol/ns110/ns53/ns224/ns304/networking_solutions_design_guidances_list.html

Content Switching Interoperability Goals

Transparency

The deployment scenarios covered in this chapter focus on the possible interoperability modes between a Catalyst 6500 and the CSM. While the available features vary between deployment modes, the main difference is in the support of the server default gateway. Offering multiple deployment modes expands the possibility of creating a viable solution to meet a variety of network design requirements. The common goal with each of these modes is providing an effective means for deploying and provisioning content switching services, while requiring little or no change to the existing data center network infrastructure.

The three deployment modes discussed in this design chapter include:

- Bridge mode
- Secure router mode
- One-arm mode

Scalability

Content switching deployments must be scalable. The logical and physical scalability must be able to increase, while requiring little or no change to the existing network infrastructure. The following factors limit scalability on a CSM:

- Connections per second (CPS)
- Concurrent connections (CC)
- Packets per second (PPS)
- Total number of VIPs

Traffic flow statistics should be gathered prior to deploying a content switch. This ensures the content switch can maintain the necessary performance numbers for CPS, CC, and PPS. These metrics are important because they determine how many connections the content switch is capable of setting up and the amount of connections it can maintain at any given point in time. These performance characteristics are further explained in an upcoming section.

High Availability

The introduction of content switching services must not compromise existing network infrastructure high availability (HA). HA for content switching deployments is broken down into two main categories: HA for the server default gateway and HA for the content switch itself. While the options vary between deployment modes, the requirement for each remains constant. You can deploy HA for the CSM in an active-standby configuration. The active-standby configuration assists in maintaining a predictable traffic flow and decreases the chance for over subscription to occur upon a failure. For example, suppose two content switches were deployed in an active-active mode and each was running at 80% of its overall capacity. Upon failure of one of the content switches, the second content switch would not be able to maintain 80% of both its capacity as well as the capacity of the failed content switch, the result would be traffic loss.

HA is also available for applications services by using stateful failover. Stateful failover retains all connection session information between devices. The end-user experiences no “disconnect” upon failure. Maintaining stateful failover for session persistence information in an e-commerce environment

is also important. The loss of session information in this environment results in the loss of transaction information like that associated with a user's shopping cart. Due to the nature of web applications, end-users can directly perceive any loss of session information involved in that session

Performance

Performance for a content switch is measured through three metrics: CPS, CC, and PPS.

CPS defines the number of connections a content switch is capable of setting up on a per second basis. This is typically limited by the hardware capabilities of the switch. CPS performance varies based on the depth of packet lookup. The further into the packet the content switch must read, the slower the CPS setup rate. Hence, connection setup rates for Layer 4 rule are much faster than for a Layer 7 rule.

CC defines the total amount of active connections that can be open on a content switch at any given time. Once the content switch reaches the limit of CCs, the new connection requests are dropped. Although the CC number is high (1 million in the CSM's case) a combination of high connection setup rates and long-lived flows can make the CC limit reachable. The CSM idle timer can be configured to age out old existing connections when they reach their time limit in order to make room for newer connections.

PPS defines the total amount of throughput capabilities for the content switch. Packet size has a direct effect on throughput capabilities. The capability of a content switch to perform hardware switching at Layer 2 is a requirement to achieve acceptable throughput performance results.

How the MSFC Communicates with the CSM

The MSFC forwards inbound traffic to the CSM through a 4-Gigabit trunk logically configured on the Catalyst 6500 backplane. To ensure traffic is sent to the correct VLAN, all traffic traversing the trunk has an associated 802.1q tag. Once traffic leaves the Catalyst 6500, these VLAN tags are stripped from the packet as depicted in Figure 3-2.

Figure 3-2 Logical View of the CSM

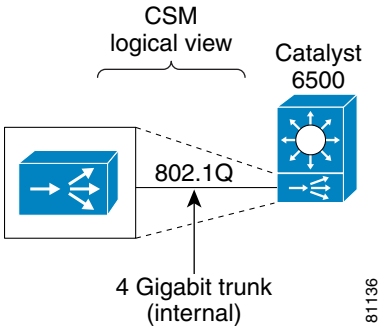

Configure the MSFC to belong to either the client-side or server-side VLANs of the CSM, but never both. Configuring the MSFC to belong to both the client and server side VLANs can cause traffic to bypass the CSM.

CSM Deployment

This section discusses two ideal locations for deploying CSMs within the enterprise data center; the data center aggregation switches and service switches.

Aggregation Switches

To ensure that traffic is properly routed through the CSM, place the CSM in a Catalyst 6500 that is in a direct path of traffic flows. The ideal location for this is in the data center aggregation switches. By placing the CSMs in the data center aggregation switches, they reside in a direct path for all traffic flows in and out of the data center. Management tasks also become easier because all Layer 2 - 3 and 4 - 7 services are housed in a central location. Additional content networking appliances, such as content engines, SSL offloaders, and content transformation devices; are directly connected to these aggregation switches as shown in Figure 3-3.

Figure 3-3 CSMs Deployed in the Data Center Aggregation Switches

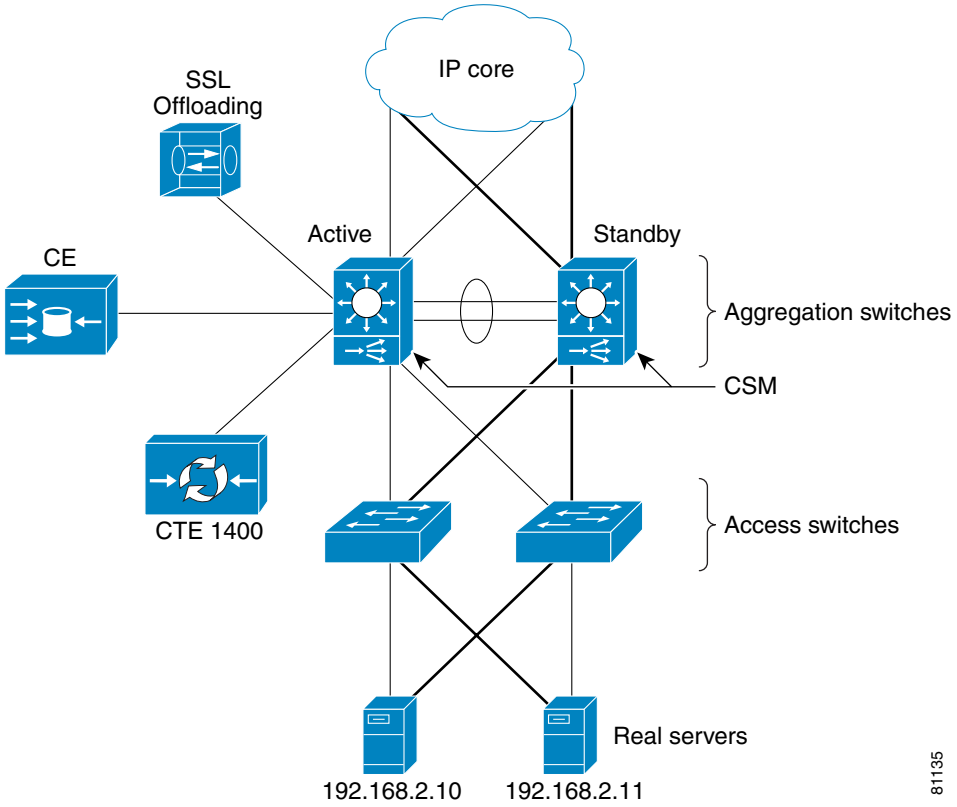

Packet Flow

Inbound

When the MSFC receives the inbound client request, it forwards the request to the CSM based on the destination IP address. The MSFC forwards all packets with a VIP as the destination IP address to the CSM. Each packet has an 802.1q tag associated with it for the incoming VLAN. This allows the CSM

to associate all packets arriving on this VLAN with the corresponding client-side VLAN. These packets are then forwarded to the server VLAN and sent back over the internal trunk to the appropriate real server.

Outbound

The return traffic from the server is considered part of the same flow. Traffic from the server is set with a destination MAC of the CSM and a destination IP address of either the client or the server-side CSM VLAN, depending on the deployment mode used.

Inbound and outbound traffics flows are illustrated in Figure 3-4.

Figure 3-4 Inbound and Outbound Traffic Flows through the CSM

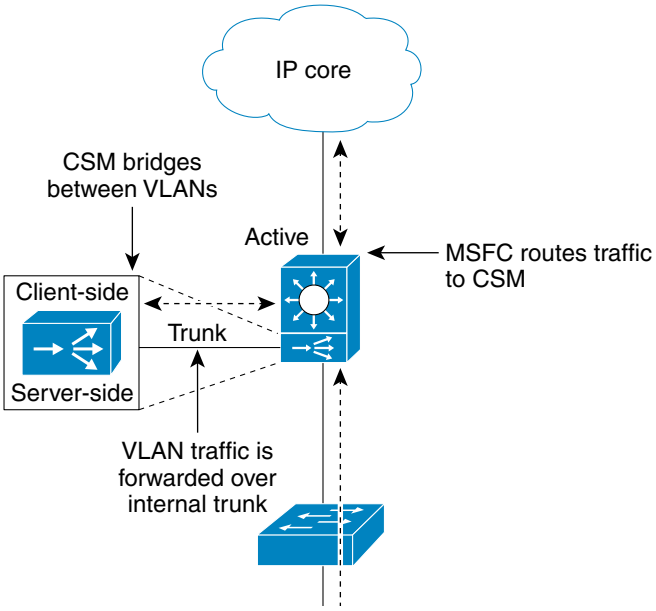

Deployment Modes

The CSM provides content switching services through a variety of deployment modes. Each of these modes is differentiated through the location of the server default gateway. Each also supports different features and functionality, providing a variety of means for meeting most network design requirements. This section covers the theoretical concepts associated with each mode. The implementation section of this design chapter provides the details of how to deploy each mode.

Bridge Mode

In bridge mode, the CSM bridges traffic flows between the client-side and server-side VLANs, rewriting the destination MAC address and destination IP address. In this deployment mode, the Catalyst 6500 MSFC provides the default gateway for the servers. This mode also allows you to leverage Cisco IOS high availability features, like Cisco's hot standby router protocol (HSRP), to support redundancy for the server default gateways. The alias command is also used on the CSM client-side VLAN for high availability. In this case, this feature provides the MSFC with a redundant next-hop address. It is also

important to utilize the VLAN keyword under the vserver configuration mode while running bridge mode. Using the VLAN keyword specifies which VLANs you can forward. The following diagram (Figure 3-5) displays a logical overview of the CSM operating in bridge mode.

Figure 3-5 Bridge Mode

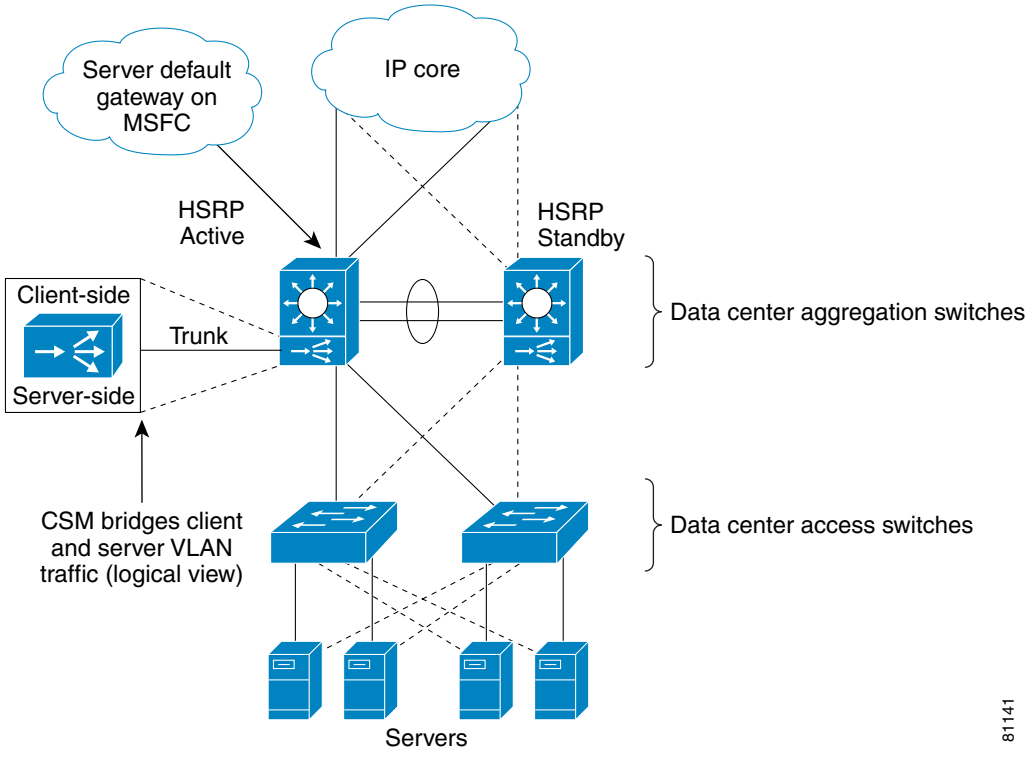

Secure Router Mode

In this mode, the server's default gateway changes location, moving from the MSFC to the CSM. Support for high availability for the default gateway is through the alias command on the CSM. The alias command provides functionality very similar to HSRP by supplying a “floating” IP and a virtual MAC address which servers point to as the default gateway. The following figure represents a logical overview of the CSM secure router mode.

Figure 3-6 CSM in Secure Router Mode

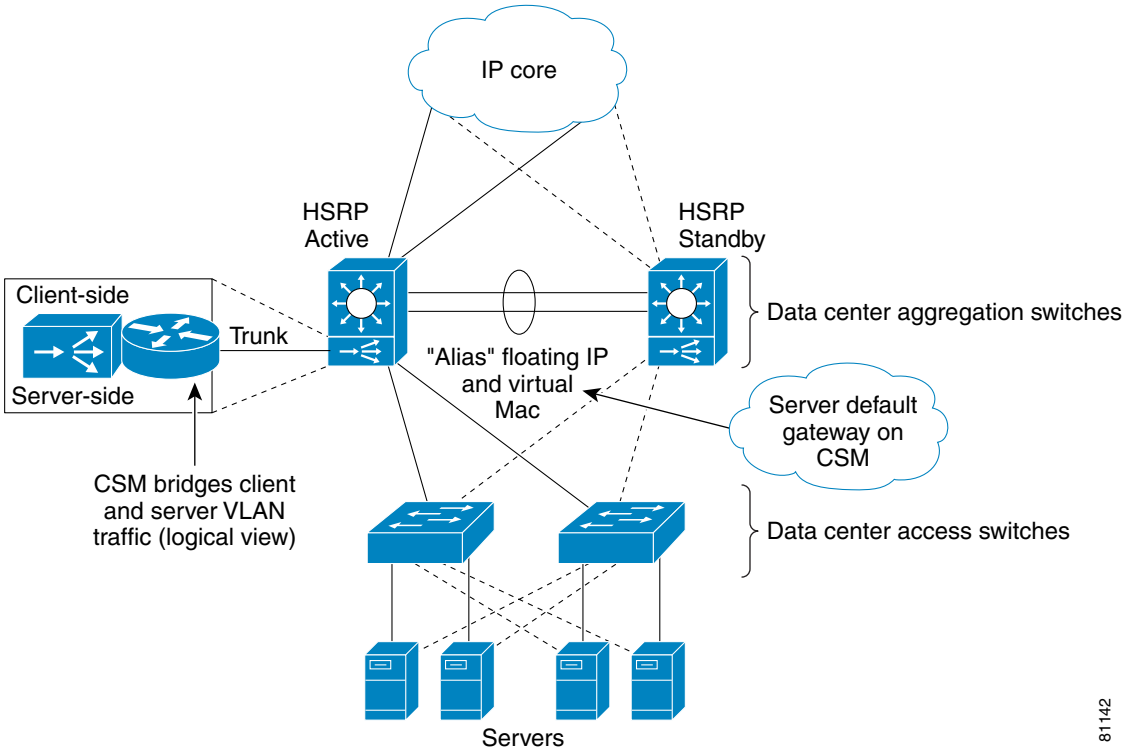

One Arm Mode

The CSM one arm mode provides a means for optimizing backend server-to-server communication. Often, servers residing in the data center need to communicate with each other or with databases residing within the data center. These communications do not necessarily need to be sent to the content switch for load balancing. This deployment mode utilizes policy based routing (PBR) to provide a means for configuring the Catalyst 6500 to only route certain traffic flows to the content switch, therefore alleviating unnecessary traffic flows from being forwarded to the content switch. This mode has other advantages as well. It allows you to configure the server default gateway on the MSFC, allowing high availability to be provided by HSRP. It also allows you to use certain Cisco IOS features, such as Private VLANs, for the server farm. A detailed explanation, along with configuration examples, is located in the Data Center Networking: Infrastructure Design Architectures SRND located at:

http://www.cisco.com/en/US/netsol/ns110/ns53/ns224/ns304/networking_solutions_design_guidances_list.html

Server CSM MSFC Communication

You can configure the servers' default gateway to be either the MSFC or the CSM. For the purpose of this design guide, there is no preferred mode. In the configurations shown, the CSM operates in bridge mode, which means it bridges the server VLANs with the client VLANs. The advantage of this is that the MSFC performs the routing functions between the server VLANs. Server-to-server traffic for separate segments (such as from 10.20.5.x to 10.20.6.x) flows all the way to the MSFC and back to the

CSM from the 10.20.6.x VLAN interface of the MSFC. The traffic from the 10.20.5.x servers going to the 10.20.6.x servers goes all the way to the MSFC and back to the CSM. The MSFC performs the routing, the CSM load balances server-to-server communication.

Figure 3-7 MSFC as the Server Default Gateway

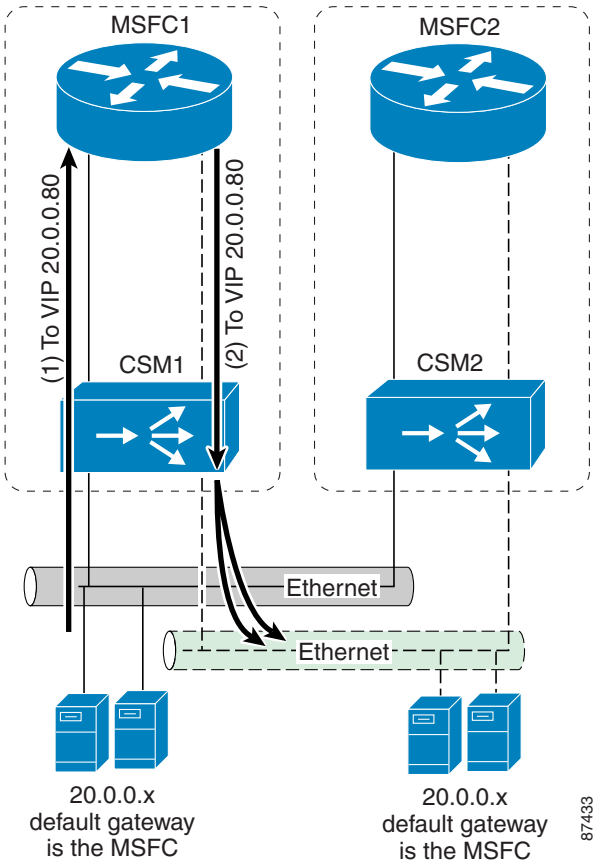

High Availability

Both the MSFC and the CSM provide several high availability options, which differ between CSM deployment modes. The MSFC uses HSRP to provide redundancy for Layer 3 interfaces that often serve as default gateways for the CSM client-side VLAN and real servers. The CSM uses a fault tolerant (FT) VLAN configuration to provide redundancy between each pair of CSM modules residing in either in the same or separate chassis.

Because the high availability configuration for the Catalyst switch and CSM is done separately, it is possible to have a scenario where an active Catalyst switch is housing a standby CSM. To maintain an easier environment for managing both physical devices and traffic flows, Cisco recommends that all active Layer 4-7 devices are connected to the active Catalyst switch.

The high availability options for each deployment mode are covered in detail in the implementation section of this chapter.

NAT (Network Address Translation)

Client NAT

Client NAT on the CSM operates in a very similar fashion to the client NAT on Cisco routers. All incoming client addresses are translated to an IP address that is part of a pool of internal addresses. This pool usually consists of IP addresses belonging to the RFC 1918 private address space.

Server NAT

Configure server NAT on the CSM either through the use of a one-to-one static map or by using port address translation (PAT) to overload the Vserver IP address (VIP). When server NAT is enabled on the CSM, clients are unable to reach the real server's IP address. The following diagram displays traffic flows with client and server NAT enabled on the CSM.

Figure 3-8 Traffic Flows with Client and Server NAT Enabled

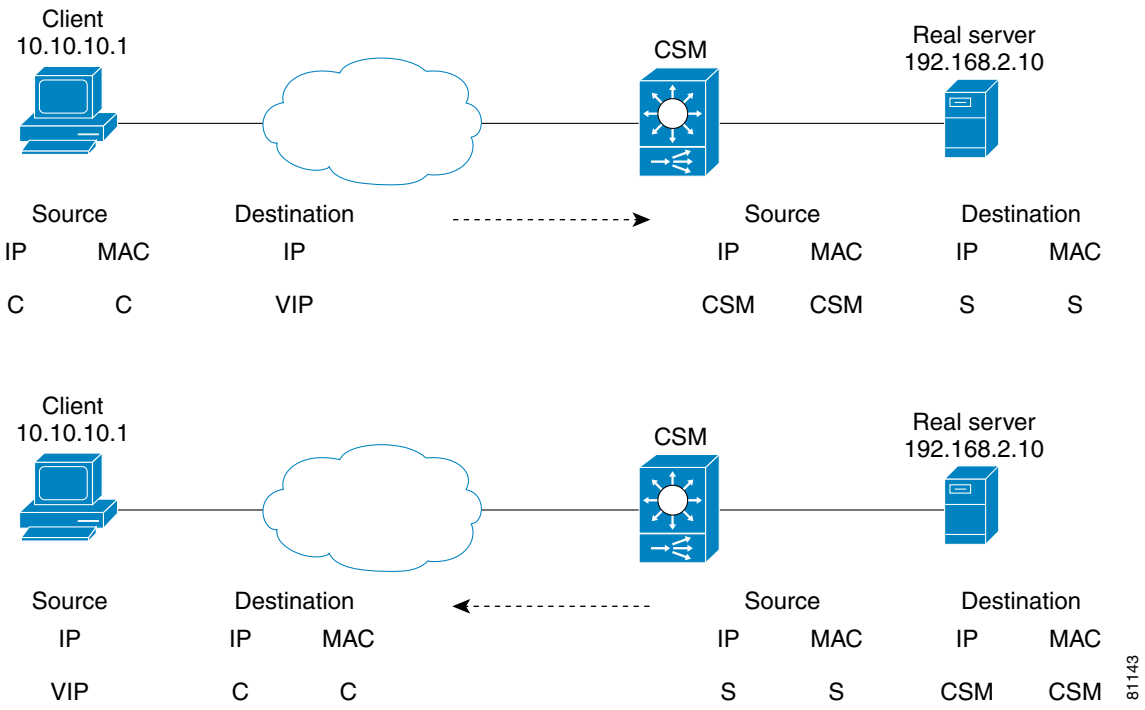

Recommendations

The recommended deployment mode must be based directly on the requirements for the network. Because these requirements vary widely from one network to another, it is not practical to recommend one deployment mode for all scenarios. Each deployment mode has its own set of requirements and caveats. The following implementation section covers the configuration details and caveats associated with each of these scenarios. Use this section to gather the information regarding each scenario to assist you in making the proper deployment mode selection.

The following diagram is a reference for the high availability discussion related to the CSM.

Figure 3-9 CSM Implementation Details

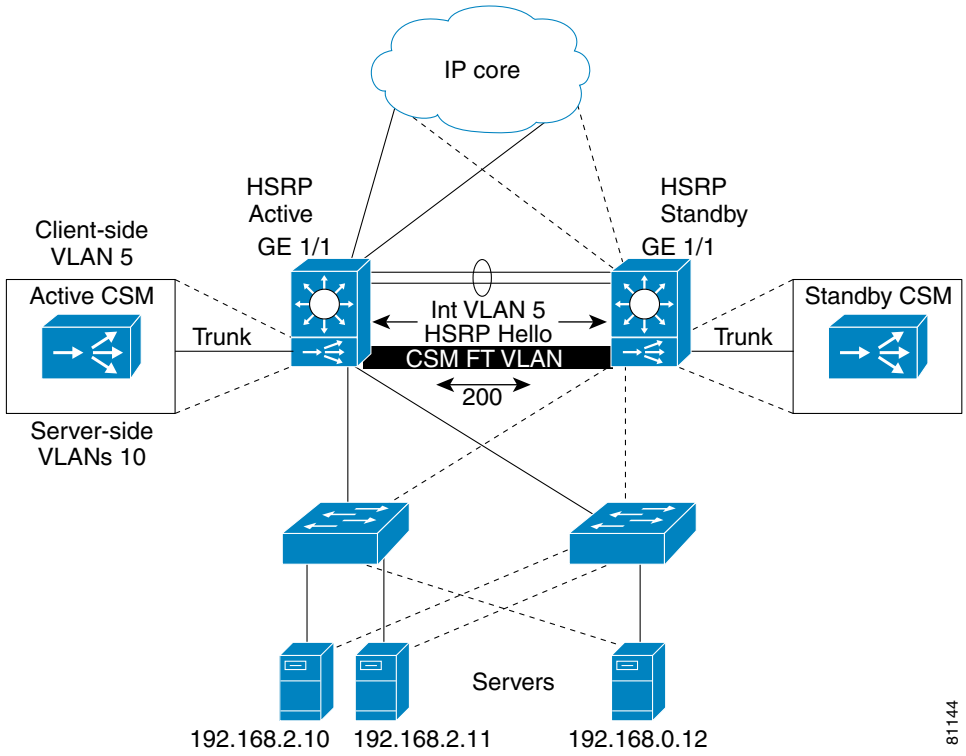

CSM High Availability

High availability for redundant CSM modules is provided through a fault tolerant (FT) VLAN configuration. When the CSMs are configured for high availability, they act as an active-standby pair. Current software releases do not allow the CSM to use the FT VLAN configuration in an active-active setup. Each pair of CSMs must be configured with a separate FT VLAN.

The CSM with higher priority assigned for the FT group becomes active, while the CSM with the lower priority takes on a standby role. By default, the hello interval of this FT information between CSMs is set to one second, and the default failover time (dead timer) is set to three seconds. Use the **show mod csm x ft** command to verify assigned VLAN, current state, priority, and timer information. The following displays the FT information for both the active and standby CSMs.

```
switch1#sh mod csm 4 ft
FT group 1, vlan 200
This box is active
 priority 30, heartbeat 1, failover 3, preemption is off
switch2#sh mod csm 4 ft
FT group 1, vlan 200
 This box is in standby state
 priority 20, heartbeat 1, failover 3, preemption is off
```

You can exchange CSM FT information over either a dedicated trunk or a Layer 2 EtherChannel. If you use a trunk to exchange the FT information, Cisco recommends that only the FT VLAN be allowed on the trunk. This is to ensure that other traffic does not overrun the link causing the loss of FT hello packets. If you use a Layer 2 EtherChannel, Cisco recommends that you enable QoS on the link guaranteeing priority to the FT hello packets. This is to ensure that a delay is not experienced with the exchange of hello packets causing both CSMs to become active. By default QoS (COS 7) is enabled on all Supervisor and MSFC combinations except for the Supervisor 2 MSFC 2 combination. If you are using this combination, you must manually enable QoS.

You can enable stateful failover on the CSM for both session and sticky information. When you enable stateful replication, state table information is replicated to the standby CSM. This ensures the standby CSM is aware of all current state information if a failure occurs. To enable replication of both state and sticky (session persistence) information perform the following.

```
switch1(config-module-csm)#vserver vs1
switch1(config-slb-vserver)#replicate csrp connection
switch1(config-slb-vserver)#replicate csrp sticky
```

To verify stateful failover is operating correctly, open a telnet connection from the client to a server through a VIP on the CSM. You can verify the current state for the active connection on the primary CSM by issuing the following command.

```
switch1#sh mod csm 4 conns
  prot vlan source               destination           state
------------------------------------------------------------------------
In TCP 5    10.15.0.6:1034        192.168.1.100:23        ESTAB
Out TCP 10   192.168.2.10:23         10.15.0.6:1034        ESTAB
```

This command establishes a telnet connection between the client and server. When the primary CSM fails, stateful replication ensures that this connection is carried over to the secondary CSM when it becomes active.

The following example displays the CSM on switch 2 is now the active CSM for FT group 1.

```
switch2#sh mod csm 4 ft
FT group 1, vlan 200
 This box is active
 priority 20, heartbeat 1, failover 3, preemption is off
```

The telnet connection was also maintained and uninterrupted between the client and server as verified in the example below.

```
switch2#sh mod csm 4 conns
    prot vlan source               destination           state
------------------------------------------------------------------------
In  TCP  6    10.15.0.20:1070       192.168.1.100:23        ESTAB
Out TCP  10   192.168.2.10:23         10.15.0.20:1070       ESTAB
```

The total length of time for the failover to occur was approximately 3.5 seconds. It should be noted that this failover time was achieved on a CSM with a very light traffic load. On a CSM with heavier traffic loads, it is possible for failover times to become longer. The results of heavy traffic loads have not been tested thoroughly by ESE. The tests are scheduled for the near future with the results published in a future revision of this chapter.

Multi-Tier Server Farm Integration

The multi-tier server farm consists of physically separate layers or tiers. These tiers are broken down into three areas: web or presentation tier, application tier, and database tier. Communication and traffic flow occurs both within each tier as well as between tiers. To provide security and manageability, each tier can be placed in a separate VLAN. This type of architecture is represented in Figure 3-10.

Figure 3-10 CSM Deployment in a Multi-Tier Server Farm

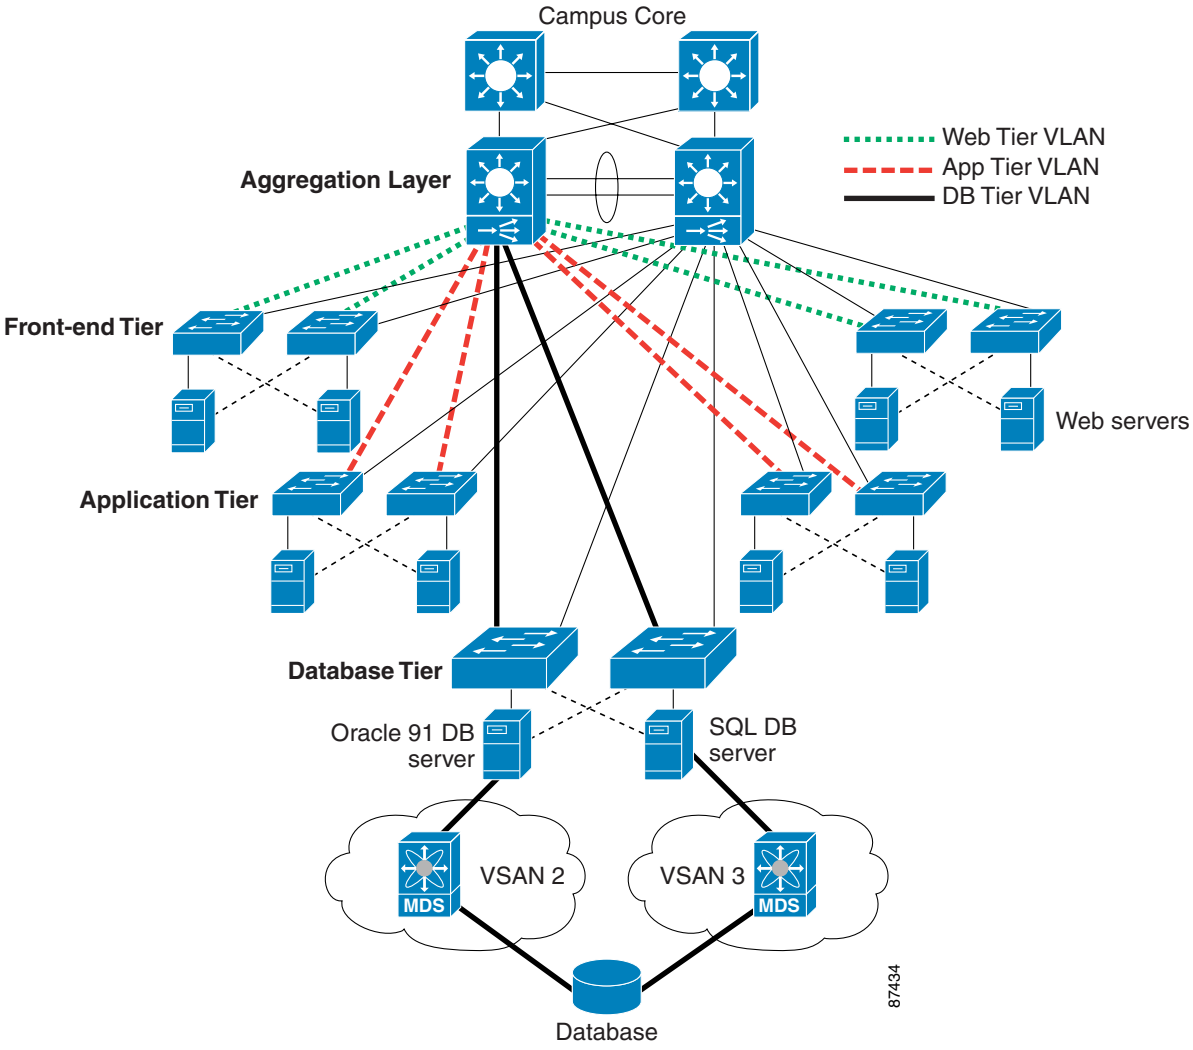

When the CSM is introduced in a multi-tier server farm it is recommended to be deployed in bridge mode. In this case, the CSM is configured with each tier's VLAN and all servers use the MSFC as the default gateway. Access control lists (ACLs) can be deployed on the MSFC to limit server farm communications. This configuration allows for easy scalability and allows the use of a single CSM to be shared across the entire server farm. In order to keep the CSM from forwarding traffic between tiers and bypassing the MSFC, the VLAN keyword can be used under each vserver configuration. This ensures that only traffic from the specified VLAN or VLANs is forwarded by the vserver. Figure 3-11 shows a detailed example of the CSM supporting multiple server VLANs while configured in bridge mode.

Figure 3-11 Example of CSM bridging multiple server VLANs

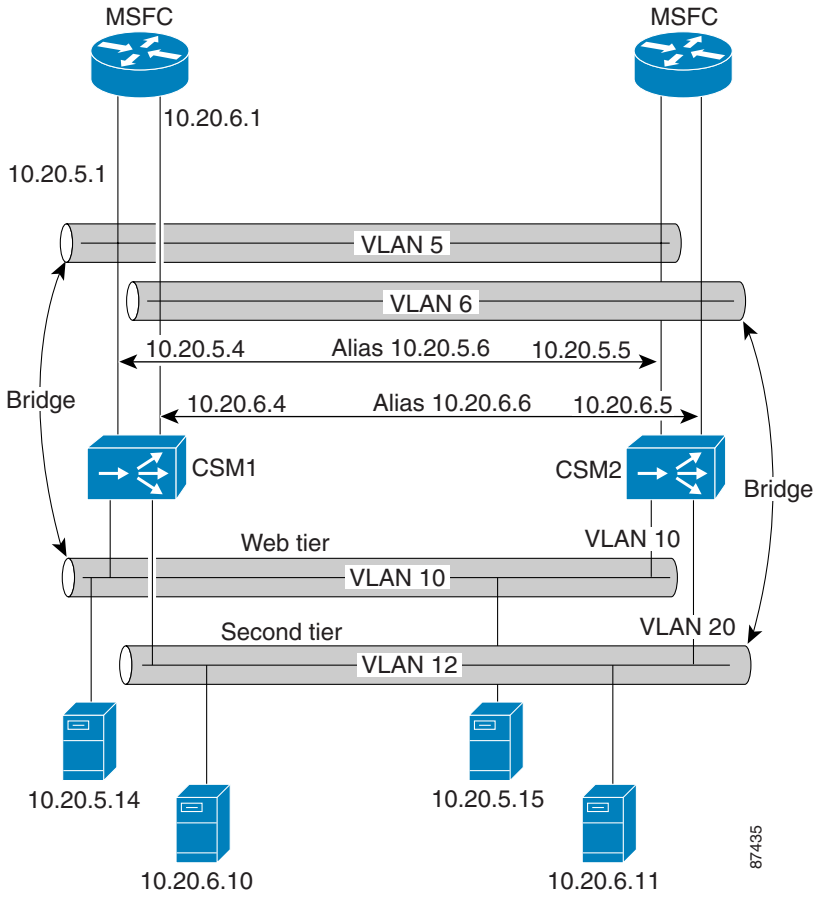

This configuration also allows for seemless integration of the Firewall Services Module (FWSM). In this environment, the FWSM would be logically placed between the MSFC and the CSM. All servers use the FWSM as their default gateway. The FWSM is configured to provide filtering both to and from and between each server farm tier.

For additional information, see Chapter 2, "Integrating the Firewall Service Module."

The CSM provides a means for deploying high performing and reliable content switching services to the Catalyst 6500 and 7600 series platforms. A variety of implementation options are available when deploying the CSM. Use the network requirements and caveats discussed in this chapter to assist you when making the decision on which mode to use when deploying the CSM and where to place it in the data center infrastructure. Refer to the ESE web page for additional data center information regarding the deployment of caching services, video distribution services, global server load balancing (disaster recovery and business continuance), and data center infrastructure design.

Integrating the Content Switching and SSL Services Modules

The SSL Service Module (SSLSM) is a service module in the Catalyst 6500 that provides offloading of Secure Socket Layer (SSL) decryption. SSL is the industry-standard method of protecting web communication using digitally encrypting data technology. The SSL protocol provides data encryption, server authentication, message integrity, and may also provide optional client side authentication. The SSL encryption engine uses a digital certificates to generate a session key. Current certificates can be either 40-bit or 128-bits in length. During SSL's initial transaction, the key initiation or handshake is the most intensive operation in SSL processing, and the most expensive operation in the handshake is the RSA private key decryption. With deployment of the SSLSM, these operations (such as, RSA private key decryption) are offloaded onto the SSLSM. SSL decryption on an SSLSM can be combined with a load balancer to provide several benefits among which are:

- Offloading SSL decryption from the servers
- HTTP session persistence across clear text and encrypted traffic
- Use of a centralized device to manage certificates

Terminology

In this chapter, a Layer 3 VLAN means a VLAN that is not trunked to the access switches and is mainly used for communication between routing devices. A Layer 3 VLAN is carried on a single trunk in the network topology, specifically the trunk + channel that runs between the two aggregation switches.

A switched VLAN interface (SVI) is a VLAN interface defined on the MSFC. A VLAN configured on the Catalyst becomes an SVI when you use the interface vlan <vlan number> command to assign it an IP address. The creation of a VLAN by itself by the command “(config) vlan <vlan number>” does not create an SVI.

In the drawings that follow, the white box that contains the FWSM, the MSFC, and the load balancer represents a Catalyst 6500, and each component is basically a blade or a daughter card in the switch.

Overview

In the data center, the SSL module is always deployed in conjunction with a load balancing device such as the Content Switching Module (CSM). This is a requirement to achieve website redundancy and scalability. The Catalyst 6500 can support up to 4 SSLSMs per chassis. If one SSL module fails, the CSM

proactively detects the failure and sends new incoming connections to the remaining SSL modules. The load balancer also provides the intelligence of routing clear text traffic to the server farm and encrypted traffic to the SSLSM.

Traffic Path

Figure 4-1 identifies the traffic path when deploying the SSLSM in the data center. Clear text traffic, such as regular HTTP GETs, goes to the CSM and the CSM distributes the requests to the servers listening on port 80. The CSM is also used to intercept encrypted HTTP traffic and forwards this traffic to the SSLSM.

The SSLSM returns the decrypted traffic to the CSM for load balancing. Because this traffic is clear-text, the CSM keeps persistence between HTTPS and HTTP.

Figure 4-1 Traffic Paths for Clear Text, Encrypted and Decrypted Traffic

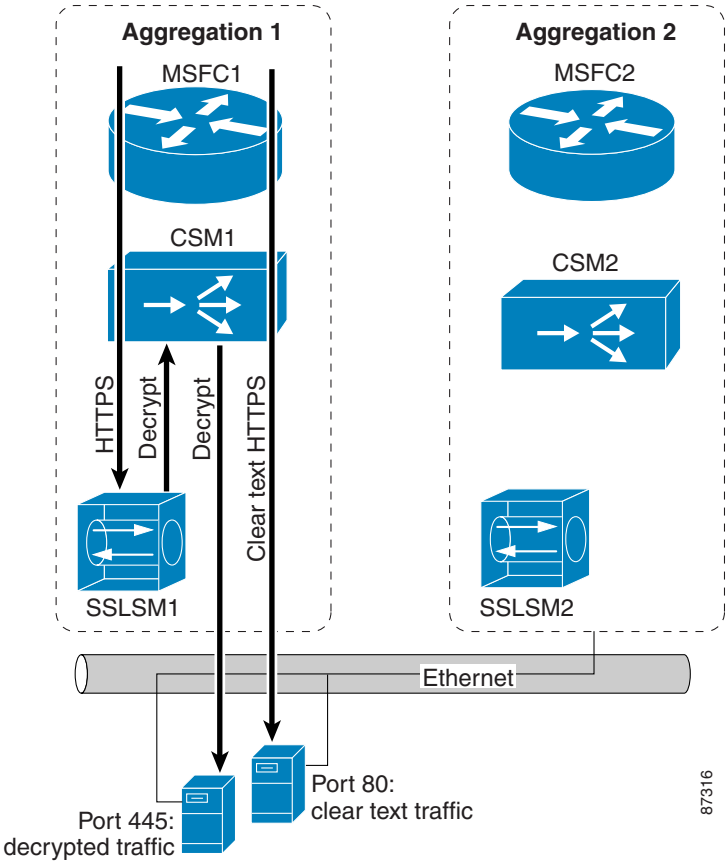

CSM SSL Communication

You can achieve communication between the CSM and the SSLSM by using one or multiple VLANs. Figure 4-2 shows VLAN allocation communications between the CSM and the SSL module: you can have either one VLAN for each domain hosted on the SSL module, or you can have a single VLAN with multiple multiplexed domains. For the purpose of this design, the choice is for the second solution. The reason for using a single VLAN is simplicity and scalability.

Figure 4-2 Connectivity Between the SSL Module and the CSM

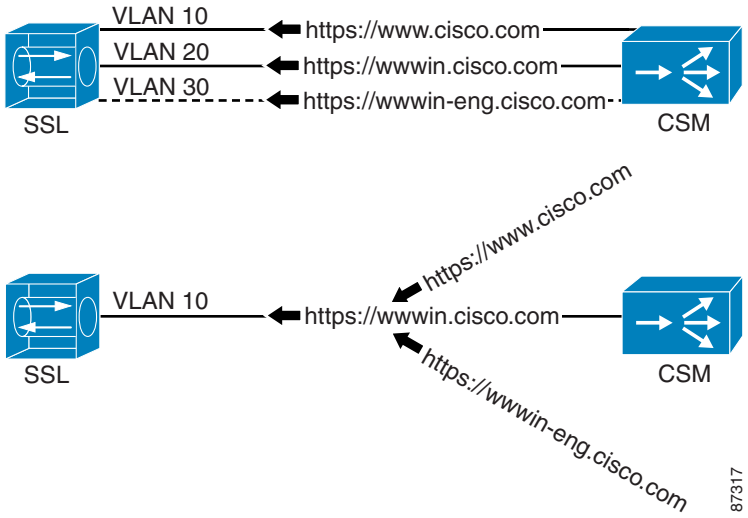

SSL MSFC communication

You can use the CSM as the default gateway for the SSL module or you can use the MSFC as the default gateway (with the CSM in the path). Figure 4-3 shows the two possible configurations: the CSM can either bridge or route between the MSFC and the SSLSM. Cisco recommends using the CSM in routed mode between the SSLSM and the MSFC. In which case, the SSL module is unaware of the MSFC and only routes traffic to the CSM.

Figure 4-3 CSM Can Bridge or Route Between the SSL and the MSFC

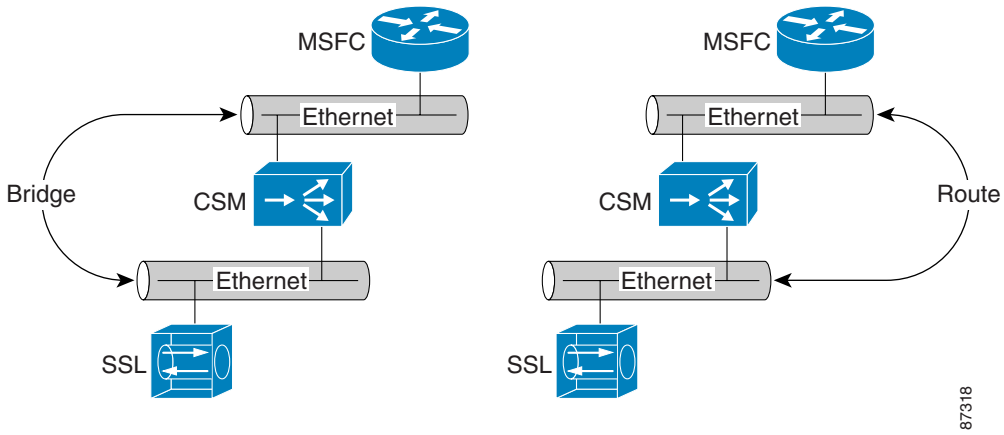

SERVERS CSM MSFC Communication

You can configure the servers' default gateway to be either the MSFC or the CSM. For the purpose of this chapter, there is no preferred mode. In the configurations shown, the CSM operates in bridge mode, which means it bridges the server VLANs with the client VLANs. The advantage of bridging is that the MSFC performs the routing functions between the server VLANs. Server-to-server traffic for separate segments (such as from 10.20.5.x to 10.20.6.x) flows all the way to the MSFC and back to the CSM from the 10.20.6.x VLAN interface of the MSFC. You can see this approach in Figure 4-4. The traffic from the 10.20.5.x servers going to the 10.20.6.x servers goes all the way to the MSFC and back to the CSM. The MSFC performs the routing, the CSM load balances server-to-server communication.

Figure 4-4 MSFC as the Server Default Gateway

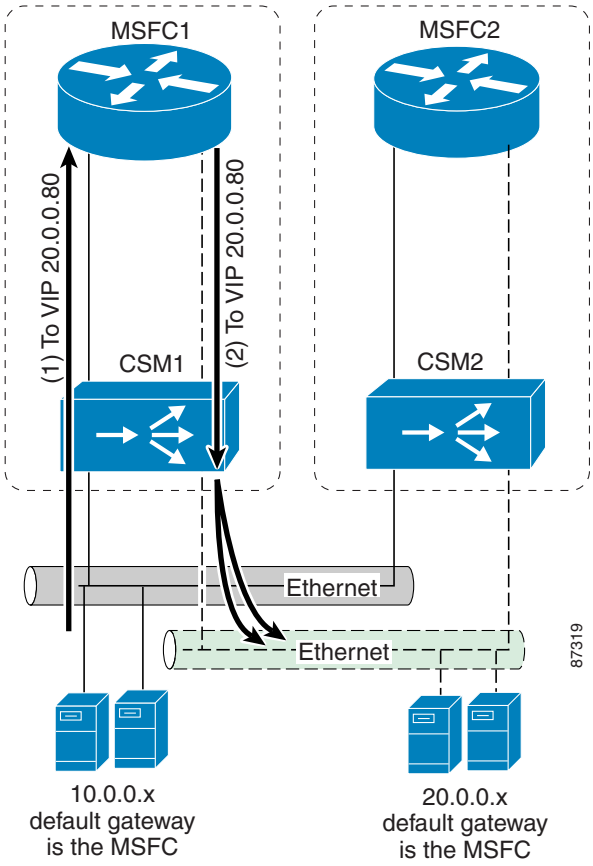

Redundancy

Use the CSM to achieve SSLSM redundancy. The CSM provides load distribution to a number of active SSLSMs.

In Figure 4-5, the SSL farm spreads across two Cisco 6500s. The CSM actively monitors the SSL modules with TCP probes on the SSL port. You can also use ICMP probes. Cisco recommends using TCP probes because TCP probes provide better health checking for the SSL module. ICMP pings succeed regardless of the certificate configuration. Therefore, a misconfigured SSL module would be perceived to be healthy with an ICMP ping. A misconfigured SSL module only answers TCP handshakes when the certificates are properly installed.

Figure 4-5 Load Distribution to the SSL Farm

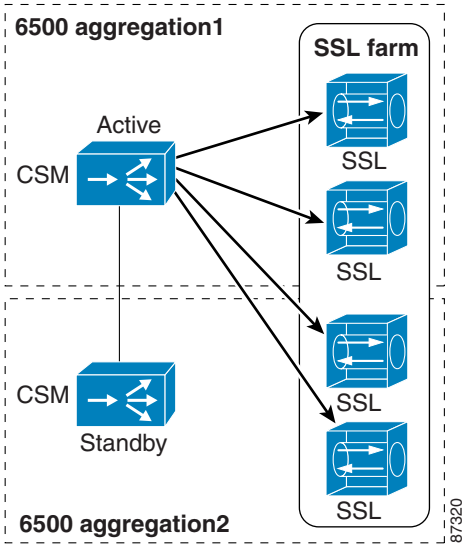

Security

When deploying SSL decryption, configure the servers to accept decrypted traffic on any port other than port 80. You must configure the HTTP servers to accept requests to port 445 or 446 and map this to the directory where you store the HTTPS page: with Apache you would change the Listen directive of the HTTP daemon that you launched for SSL traffic.

A practical example of mapping between ports and domain names is using port 445 for the decrypted traffic of https://www.cisco.com or 446 for the decrypted traffic of https://wwwin.cisco.com.

It is important to prevent client HTTP traffic from going to the server port 445 or 446. You should allow only traffic coming from the SSL offloader to the HTTP server on port 445 and 446.

The CSM provides this protection by only allowing the traffic from the SSL module (SSL VLAN specifically) to get access to the servers via the virtual IP address. Access lists on the MSFC or on a firewall should also drop client traffic destined to the ports for the decrypted traffic (445, 446).

Scalability

The scalability numbers for the SSL module are as follows:

- 3K RSA/s with no session resumption (1024-bit RSA key)
- 3.9K RSA/s with session resumption (1024-bit RSA key)
- 300 Mbps throughput with RC4 and MD5 (symmetric)
- 60K concurrent sessions (64k SSL “connections” to the clients + 64k HTTP “connections” to the servers)
- 1 Gbps connection full duplex to the 6500 backplane
- 256 proxy servers

- 356 key pairs
- 356 certificates

As a result of these numbers you can expect each CSM to load balance a maximum of 10 to 15 SSL modules. These numbers are given by throughput, or Layer 5 setup rate ratio of the two modules.

Data Center Configurations Description

Topology

Figure 4-6 shows the topology used in this design guide. The MSFC routes traffic from the core network to the CSM modules. The MSFC advertises the 10.20.5.x subnet. The VIPs belong to this subnet, for example 10.20.5.80 or 10.20.5.90. The SSL modules reside on the 10.20.3.x subnet. The SSL modules' default gateway is the CSM (the CSM alias to be precise). Figure 4-6 shows their IP addresses.

Figure 4-6 CSM + SSL Layer 3 Tested Topology

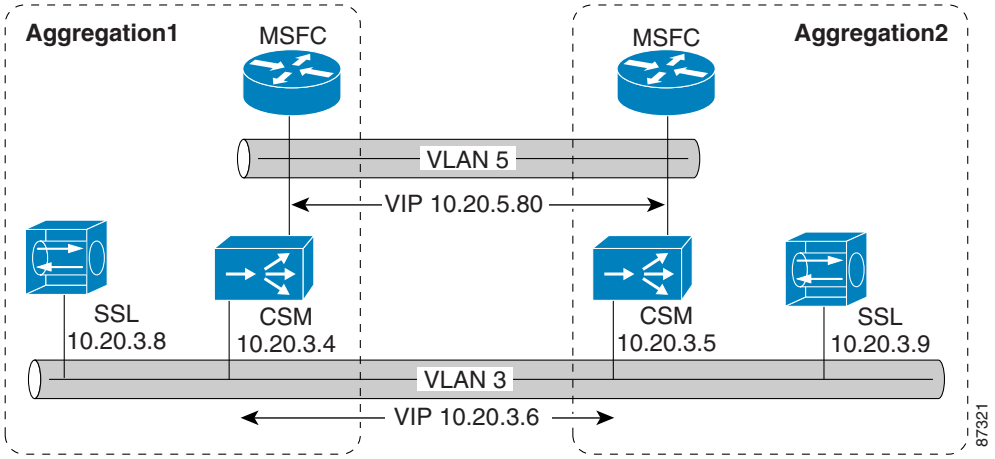

Figure 4-6 also shows the Layer 3 design, which is the IP addresses used for “routing”. This graphic does not show the Layer 4 configuration which is the set of IP addresses used to define the “services”. When the CSM receives a request for “https://www.cisco.com”, the CSM does not forward this to the “Layer 3" IP of the SSL, it sends it to the “service” IP of the SSL module. Similarly, when the SSL module sends traffic back to the CSM, it does not send it to the alias IP of the CSM, it sends it to the “service” IP of the CSM (the vserver).

Figure 4-7 CSM + SSL Layer 4 Topology (Service IP Addresses)

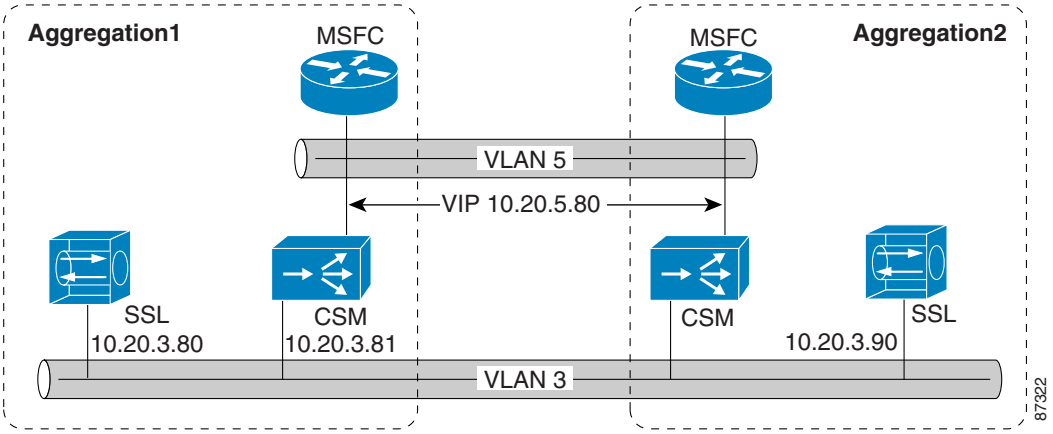

Figure 4-7 shows the topology with the service IP addresses:

- SSL module 1 offers services to the CSM on the 10.20.3.80 IP
- SSL module 2 offers services to the CSM on the 10.20.3.90 IP
- The CSM offers services to the SSL module on the 10.20.3.81 IP

Suppose you need to configure and decrypt HTTPS for multiple VIPs (each mapping to a different domain) such as, 10.20.5.80 to 10.20.5.90. The SSL module receiving traffic with destination IP equal to 10.20.3.80 must differentiate between the domains. This is achieved by rewriting the destination port.

As you know, HTTPS uses port 443 (the SSL port). The CSM intercepts this traffic and changes the destination port to help the SSL module differentiate traffic for the domains.

For example:

- VIP1: https://www.cisco.com (10.20.5.80) maps to port 445
- VIP2: https://wwwin.cisco.com (10.20.5.90) maps to port 446

As a result, the following mappings are on SSL module 1:

- VIP1: https://www.cisco.com - 10.20.3.80 port 445
- VIP2: https://wwwin.cisco.com - 10.20.3.80 port 446

The following mappings exist on SSL module 2:

- VIP1: https://www.cisco.com - 10.20.3.90 port 445
- VIP2: https://wwwin.cisco.com - 10.20.3.90 port 446

Traffic decrypted by the SSL module needs to use the load balancing service of the CSM. This service is identified by the IP 10.20.3.81. As a result, the CSM has the following mappings:

- VIP1: https://www.cisco.com - 10.20.3.81 port 445
- VIP2: https://wwwin.cisco.com - 10.20.3.81 port 446

The servers and the MSFC communicate using the CSM in either of the two following modes:

- Routed mode: the CSM routes between VLANs (the CSM is the default gateway for the servers)
- Bridge Mode: the CSM bridges client and server VLANs

For the purpose of this design, both modes are equal. In this specific setup, you see configurations for bridge mode, because the MSFC is the servers' default gateway. All the server-to-server communication flows through the CSM and the MSFC. The VLAN configuration on the CSM forces the server-to-server traffic to leave the CSM, go to the MSFC, and come back to the CSM. Figure 4-8 shows this topology. The default gateway for the servers in the web-tier is the MSFC interface 10.20.5.1 and the default gateway for the servers in the 2nd tier subnet is the MSFC interface 10.20.6.1.

You can find details about the Layer 2 design of the data center infrastructure in the *Data Center Networking: Infrastructure Design Architectures* available at:

http://www.cisco.com/en/US/netsol/ns110/ns53/ns224/ns304/networking_solutions_design_guidances_list.html

Figure 4-8 CSM + MSFC and the Servers

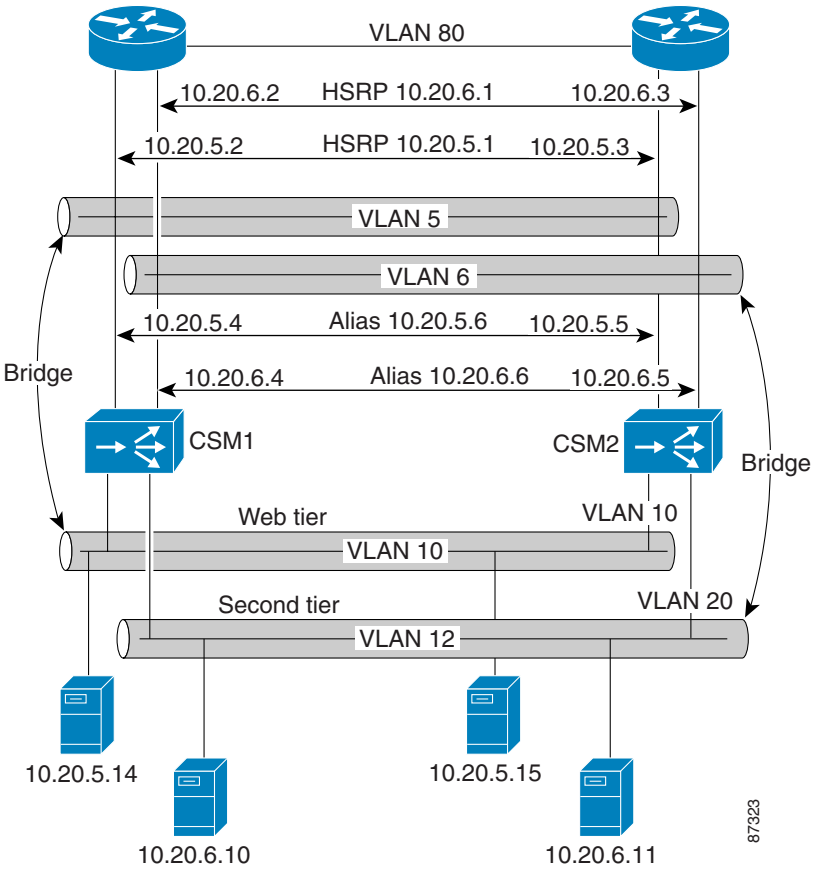

Layer 2

For the basic configuration of VLANs, channels, and spanning-tree, refer to the *Data Center Networking: Infrastructure Design Architectures SRND*. This section covers the configuration of the VLAN topology to allow communication between the CSM, SSL module and the MSFC.

Configuring VLANs on the 6500

You must create the necessary VLANs on the Catalyst 6500 first. For this configuration, those VLANS are:

- VLAN3: for the communication between the CSM and the SSLSM
- VLAN5: one of the client side VLANs of the CSM (this could be the VLAN for the web-servers)
- VLAN6: another client side VLAN of the CSM (this could be the VLAN for the application servers)
- VLAN10: merged with VLAN5 by the CSM
- VLAN12: merged with VLAN6 by the CSM
- VLAN100: FT vlan for the CSM

In addition to these VLANs, you need to configure the management VLAN for the SSLSM (VLAN 110 in the configurations).

The following commands configure the above VLANs:

```
mp_agg1(config)# vtp domain mydomain
mp_agg1(config)# vtp mode transparent
mp_agg1(config)#vlan 3
mp_agg1(config-vlan)# name SSLVLAN
mp_agg1(config)#vlan 5
mp_agg1(config-vlan)# name WEBTIERVLAN
[continue for the other VLANs]
```

Configure the VLANs on both Aggregation1 and Aggregation2. Trunk these VLANs between the two 6500s on the previously created channel:

```
interface Port-channel2
 no ip address
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 3,5,6,10,12,100
 switchport mode trunk
 spanning-tree guard loop
!
```

Note Make sure to trunk VLAN 100, the FT VLAN for the CSM.

For security reasons, enable tagging on all the VLANs:

```
vlan dot1q tag native
```

Choose the Spanning-Tree algorithm, our recommendation is Rapid PVST+:

```
spanning-tree mode rapid-pvst
```

Make Aggregation1 the root switch:

```
spanning-tree vlan 3,5,6,10,12,80,100 root
```

Make Aggregation2 the secondary root switch:

```
spanning-tree vlan 3,5,6,6,10,12,80,100 root secondary
```

Configuring VLANs on the CSM

Configure the client and server side VLANs on the CSM and bridge them. The following is the configuration for Aggregation1, the configuration on Aggregation2 is similar with the exception of the highlighted fields:

```
module ContentSwitchingModule 5
vlan 3 server
  ip address 10.20.3.4 255.255.255.0
  alias 10.20.3.6 255.255.255.0
!
 vlan 5 client
  ip address 10.20.5.4 255.255.255.0
  alias 10.20.5.6 255.255.255.0
!
vlan 6 client
  ip address 10.20.6.4 255.255.255.0
  alias 10.20.6.6 255.255.255.0
!
vlan 10 server
ip address 10.20.5.4 255.255.255.0
!
vlan 12 server
  ip address 10.20.6.4 255.255.255.0
!
ft group 1 vlan 100
  priority 10
  heartbeat-time 5
 failover 4
!
```

Notice the following key points:

- The CSM uses VLAN3 to send traffic to the SSL offloader and to receive it back.
- In this example, the servers belong to two separate broadcast domains: 10.20.5.x and 10.20.6.x. You might not need to use two, you might just need one, in which case you would only bridge VLAN 5 with VLAN 10
- Use the same IP address statement: “ip address 10.20.5.4" on both VLANs to bridge between VLAN5 and VLAN10.
- Use the same IP address statement: “ip address 10.20.6.4" to bridge between VLAN6 and VLAN12.

Configuring VLANs on the SSLSM

Configure the required VLANs on the SSLSM using the **ssl-proxy vlan** command.

This is the configuration on Aggregation1 (from the SSLSM console):

```
ssl-proxy vlan 3
 ipaddr 10.20.3.8 255.255.255.0
 gateway 10.20.3.6
```

This is the configuration on Aggregation2 (from the SSLSM console):

```
ssl-proxy vlan 3
 ipaddr 10.20.3.9 255.255.255.0
 gateway 10.20.3.6
```

Layer 3

On the Catalyst switches there are VLANs used just for bridging and VLANs that are used for routing (called SVIs). This section shows which SVIs to configure and how to configure them. The diagrams for this section are Figure 4-6 and Figure 4-8.

Configuring IP Addresses on the MSFCs

Configure MSFC SVIs only on VLAN5 and VLAN6, do not configure SVIs on VLAN10, VLAN12, VLAN3. The following is the configuration for Aggregation1.

```
interface Vlan5
 ip address 10.20.5.2 255.255.255.0
 no ip redirects
 standby 1 ip 10.20.5.1
 standby 1 priority 110
 standby 1 preempt
!
interface Vlan6
 ip address 10.20.6.2 255.255.255.0
 no ip redirects
 standby 1 ip 10.20.6.1
 standby 1 priority 110
 standby 1 preempt
!
```

The configuration on Aggregation2 is the same with the exception of the highlighted fields.

Configuring IP Addresses on the CSM

You already configured the IP address on the CSM by configuring the VLANs (see the Layer 2 section).

Configuring IP Addresses on the SSLSM

The configuration of the IP address is the same as in the Layer 2 section. Notice that the gateway statement points to the ALIAS IP address of the CSM:

```
ssl-proxy vlan 3
 ipaddr 10.20.3.8 255.255.255.0
 gateway 10.20.3.6
```

If the CSM is not reachable, the output from the “show ssl-proxy service” command has the services marked as down.

Layer 4 and 5

The Layer 4 configuration on the CSM consists in two separate vservers:

- One vserver for the encrypted traffic which is used to intercept HTTPS requests and send them to the SSL offloading device
- A second vserver for the decrypted traffic used to load balance the decrypted traffic among the available servers. The configuration of these two are purposefully split.

The other important point to notice is the IP address used in the “real” configuration of the CSM, as well as the IP address used in the service configuration of the SSL module. These are the “service” IP addresses, not the regular IP addresses assigned in the Layer 3 configuration.

CSM Configuration to Intercept HTTPS Traffic

The CSM module intercepts HTTPS requests (port 443) and distributes them to the available SSL modules. The CSM treats the SSL modules like regular servers. You use the exact same configuration on Aggregation2. Notice the port address translation to port 445 in the server farm SSLBLADE-VIP1. This is explained in further detail later.

The CSM translates the destination IP of 10.20.5.80 to the 10.20.3.80 or 10.20.3.90 IP address and also rewrites the Layer 4 port to 445. A TCP probe monitors the SSL module capability of establishing connections on port 443.

```
vserver HTTPS-VIP1
 virtual 10.20.5.80 tcp https
 vlan 5
 serverfarm SSLBLADE-VIP1
 persistent rebalance
inservice
!
serverfarm SSLBLADE-VIP1
 nat server
 no nat client
 real 10.20.3.80 445
   inservice
 real 10.20.3.90 445
   inservice
probe TCP
!
probe TCP tcp
 interval 2
 failed 3
!
```

SSLSM Configuration

Configure the SSL module to accept incoming connections on 10.20.5.80 port 445 (virtual ipaddr 10.20.3.80 protocol tcp port 445). Use the server ipaddr to forward the traffic to the CSM (notice the use of 10.20.3.81, not the use of 10.20.3.6). The certificate configuration indicates which certificate is associated with this service.

Notice that the SSL module forwards traffic to the CSM vserver 10.20.3.81 and that it sends the traffic to port 445.

```
ssl-proxy service VIP10.20.5.80
 virtual ipaddr 10.20.3.80 protocol tcp port 445
 server ipaddr 10.20.3.81 protocol tcp port 445
 certificate rsa general-purpose trustpoint www
 inservice
```

Load Balancing the Decrypted Traffic

The CSM accepts the decrypted traffic from the SSL module on VLAN3 (notice VLAN3 only - this is enforced with the “vlan 3" statement in the vserver configuration) and uses 10.20.5.81 as the IP address for decrypted traffic. Port 445 identifies the VIP which is performing the decryption. The following configuration is the same on Aggregation1 and Aggregation2.

Note The Web-server with the SSL content are configured to listen on port 445. If you do not configure them to listen to 445 (or whatever port designated for decrypted traffic) this configuration does not work.

```
vserver WEBDECRYPT-VIP1
  virtual 10.20.3.81 tcp 445
  vlan 3
  serverfarm WEB-VIP1
  persistent rebalance
  inservice
!
serverfarm WEB-VIP1
  nat server
  no nat client
  real 10.20.5.14
   inservice
  real 10.20.5.15
   inservice
  probe TCP
!
```

Returning Decrypted HTTP Responses to the SSLSM

HTTP responses that require encryption are originated by the servers. In theory, the CSM could route the returning decrypted traffic to the MSFC bypassing the SSL offloader. In fact, the destination IP address is the client IP. This does not happen because the CSM remembers the flow created by the SSL decrypted traffic and is capable of mapping the returning traffic back to its original VLAN. If the HTTP request was decrypted, the incoming VLAN of the decrypted traffic is 3 (for this design). The corresponding HTTP response goes back to VLAN 3 because the CSM remembers the flow for the HTTP request.

Security

Specify the source VLAN in the vserver configuration (in this example is VLAN 3). You do not want a client from the Internet t send HTTP traffic to the servers on port 445 or 446. A connection to these ports is supposed to be encrypted by the SSL module.

A client request for HTTP to port 445 or 446 comes in from VLAN 5. As such, it does not match the vserver WEBDECRYPT-VIP1 or WEBDECRYPT-VIP2.

```
vserver WEBDECRYPT-VIP1
  virtual 10.20.3.81 tcp 445
  vlan 3
  serverfarm WEB-VIP1
  persistent rebalance
  inservice
!
vserver WEBDECRYPT-VIP2
  virtual 10.20.3.81 tcp 446
  vlan 3
  serverfarm WEB-VIP1
  persistent rebalance
  inservice
!
```

In addition to these security measures, you could configure access-lists to block external traffic destined to 10.20.3.81 by using an ACL on the MSFC.

Multiple VIPs

The data center in this design guide hosts two different VIPs: 10.20.5.80 and 10.20.5.90. The CSM intercepts the traffic and performs destination NAT to the IP address of the SSLSM.

CSM intercepts the SSL traffic:

```
vserver HTTPS-VIP1
  virtual 10.20.5.80 tcp https
  vlan 5
  serverfarm SSLBLADE-VIP1
  persistent rebalance
inservice

vserver HTTPS-VIP2
  virtual 10.20.5.90 tcp https
  vlan 5
  serverfarm SSLBLADE-VIP2
  persistent rebalance
inservice
!
The CSM does destination port rewrite:

serverfarm SSLBLADE-VIP1
  nat server
  no nat client
  real 10.20.3.80 445
   inservice
  real 10.20.3.90 445
   inservice
probe TCP
!
serverfarm SSLBLADE-VIP2
  nat server
  no nat client
  real 10.20.3.80 446
   inservice
  real 10.20.3.90 446
   inservice
probe TCP
!
SSL decrypts the traffic for 10.20.5.80 on port 445 and for 10.20.5.90 on port 446:
ssl-proxy service VIP10.20.5.80
 virtual ipaddr 10.20.3.80 protocol tcp port 445
 server ipaddr 10.20.3.81 protocol tcp port 445
 certificate rsa general-purpose trustpoint www
 inservice

ssl-proxy service VIP10.20.5.90
 virtual ipaddr 10.20.3.80 protocol tcp port 446
 server ipaddr 10.20.3.81 protocol tcp port 446
 certificate rsa general-purpose trustpoint wwwin
 inservice

CSM load balances the decrypted traffic:
vserver WEBDECRYPT-VIP1
  virtual 10.20.3.81 tcp 445
  vlan 3
  serverfarm WEB-VIP1
  persistent rebalance
  inservice

vserver WEBDECRYPT-VIP2
  virtual 10.20.3.81 tcp 446
```

```
  vlan 3
  serverfarm WEB-VIP2
  persistent rebalance
  inservice
!
```

As you can see in the configuration, the more VIPs you host, the more Layer 4 ports you allocate. All the vservers on the CSM continue to use the same "service" IP address. This allows you to cut and paste the same configuration and just changing the Layer 4 port, you do not have to burn IP addresses.

Persistence

The key point of using an SSL module together with a CSM is to achieve session persistence for HTTP on both clear text and encrypted traffic. The following configuration shows how to achieve this:

```
vserver HTTP-VIP1
  virtual 10.20.5.80 tcp www
  vlan 5
  serverfarm WEB-VIP1
  advertise active
  sticky 10 group 1
  replicate csrp sticky
  persistent rebalance
  parse-length 4000
  inservice
!
vserver WEBDECRYPT-VIP1
  virtual 10.20.3.81 tcp 445
  vlan 3
serverfarm WEB-VIP1
  replicate csrp sticky
sticky 10 group 1
persistent rebalance
  parse-length 4000
  inservice
!
sticky 1 cookie cookie-server timeout 10
```

As you can see, the sticky group that learns the cookies from the servers is applied to both vserver HTTP-VIP1 and WEBDECRYPT-VIP1, which means that session persistence is preserved when transitioning from HTTP to HTTPS.

Configurations

The configurations that you find in this section refer to the topology that you can see in Figure 4-9. Notice that the CSM and SSL modules are not externally connected, the diagram shows them as external devices just to clarify the VLANs being used and their IP addresses.

Figure 4-9 Topology Diagram

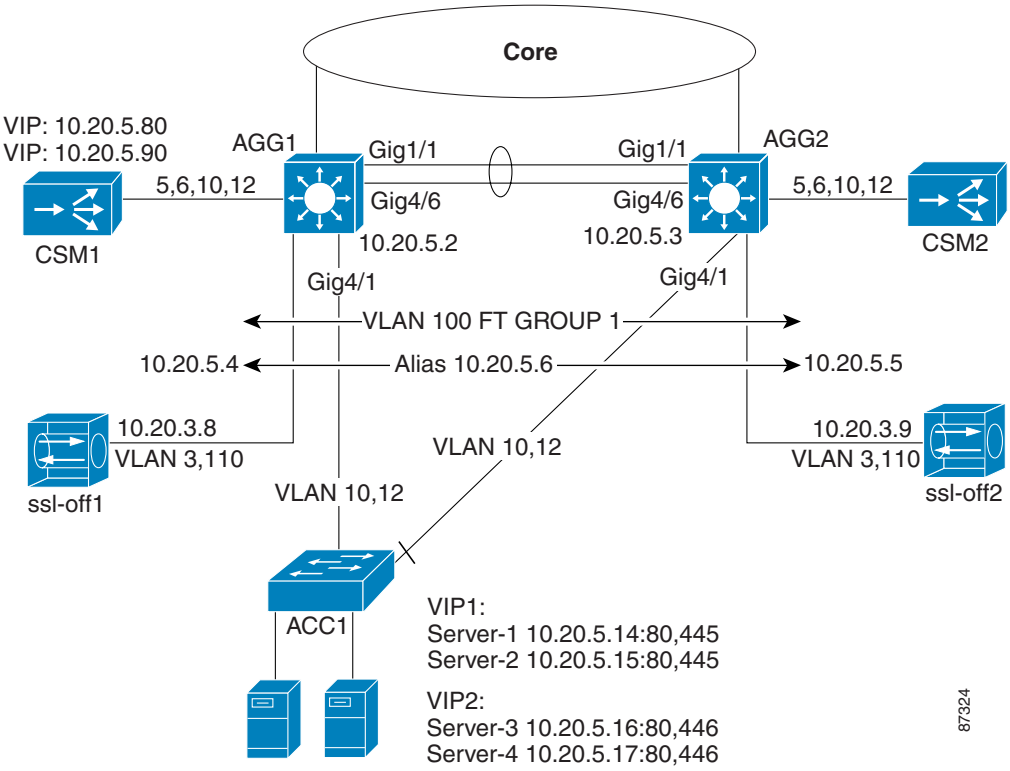

Aggregation1

```
agg1#show run
!
ssl-proxy module 3 allowed-vlan 3,110
vtp domain mydomain
vtp mode transparent
ip subnet-zero
!
!
no ip domain-lookup
!
mls flow ip destination
mls flow ipx destination
mls sampling packet-based 1024 4096
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
spanning-tree vlan 3,5,6,10,12,80,100 priority 8192
!
module ContentSwitchingModule 5
 vlan 5 client
  ip address 10.20.5.4 255.255.255.0
  alias 10.20.5.6 255.255.255.0
!
 vlan 3 server
  ip address 10.20.3.4 255.255.255.0
```

```
  alias 10.20.3.6 255.255.255.0
!
vlan 6 client
  ip address 10.20.6.4 255.255.255.0
  alias 10.20.6.6 255.255.255.0
!
 vlan 10 server
  ip address 10.20.5.4 255.255.255.0
!
 vlan 12 server
  ip address 10.20.6.4 255.255.255.0
!
probe TCP tcp
  interval 2
  failed 3
!
 probe ICMP icmp
  interval 3
  failed 5
!
serverfarm SSLBLADE-VIP1
  nat server
  no nat client
  real 10.20.3.80 445
   inservice
  real 10.20.3.90 445
   inservice
  probe TCP
!
serverfarm SSLBLADE-VIP2
  nat server
  no nat client
  real 10.20.3.80 446
   inservice
  real 10.20.3.90 446
   inservice
  probe TCP
!
 serverfarm WEB-VIP1
  nat server
  no nat client
  real 10.20.5.14
   inservice
  real 10.20.5.15
   inservice
  probe TCP
!
 serverfarm WEB-VIP2
  nat server
no nat client
  real 10.20.5.16
   inservice
  real 10.20.5.17
   inservice
  probe TCP
!
 sticky 1 cookie cookie-server timeout 10
 sticky 2 cookie cookie-server timeout 10
!
vserver HTTP-VIP1
  virtual 10.20.5.80 tcp www
  vlan 5
  serverfarm WEB-VIP1
  advertise active
```

```
  sticky 10 group 1
  replicate csrp sticky
  persistent rebalance
  parse-length 4000
  inservice
!
vserver HTTP-VIP2
  virtual 10.20.5.90 tcp www
  vlan 5
  serverfarm WEB-VIP2
  advertise active
  sticky 10 group 2
  replicate csrp sticky
  persistent rebalance
  parse-length 4000
  inservice
!
vserver HTTPS-VIP1
  virtual 10.20.5.80 tcp https
  vlan 5
  serverfarm SSLBLADE-VIP1
persistent rebalance
  inservice
!
vserver HTTPS-VIP2
  virtual 10.20.5.90 tcp https
  vlan 5
serverfarm SSLBLADE-VIP2
persistent rebalance
  inservice
!
vserver WEBDECRYPT-VIP1
  virtual 10.20.3.81 tcp 445
  vlan 3
  serverfarm WEB-VIP1
  sticky 10 group 1
  replicate csrp sticky
  persistent rebalance
  parse-length 4000
  inservice
!
vserver WEBDECRYPT-VIP2
  virtual 10.20.3.81 tcp 446
  vlan 3
serverfarm WEB-VIP2
  sticky 10 group 2
  replicate csrp sticky
  persistent rebalance
  parse-length 4000
  inservice
!
ft group 1 vlan 100
  priority 20
  heartbeat-time 5
  failover 4
!
redundancy
 mode rpr-plus
 main-cpu
  auto-sync startup-config
  auto-sync running-config
  auto-sync standard
!
vlan dot1q tag native
```

```
!
vlan 3
 name SSLVLAN
!
vlan 5
 name WEBTIERVLAN
!
vlan 6
 name APPTIERVLAN
!
vlan 10
 name WEBSERVERVLAN
!
vlan 12
 name APPSERVERVLAN
!
vlan 80
 name Layer3_VLAN
!
vlan 100
 name CSM_fault_tolerant
!
vlan 110
 name ssl_admin
!
interface Port-channel2
 no ip address
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 3,5,6,10,12,80,100,110
 switchport mode trunk
 spanning-tree guard loop
!
interface GigabitEthernet1/1
 description to_mp_agg2
 no ip address
 logging event link-status
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 3,5,6,10,12,80,100
 switchport mode trunk
 channel-group 2 mode active
 channel-protocol lacp
!
interface GigabitEthernet4/1
 description to_mp_acc1
 no ip address
 switchport
!
interface GigabitEthernet4/6
 description to_mp_agg2
 no ip address
 logging event link-status
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 3,5,6,10,12,80,100,110
 switchport mode trunk
 channel-group 2 mode active
 channel-protocol lacp
!
interface Vlan5
 ip address 10.20.5.2 255.255.255.0
 no ip redirects
 standby 1 ip 10.20.5.1
```

```
 standby 1 priority 110
 standby 1 preempt
!
interface Vlan6
 ip address 10.20.6.2 255.255.255.0
 no ip redirects
 standby 1 ip 10.20.6.1
 standby 1 priority 110
 standby 1 preempt
!
interface Vlan80
 ip address 10.80.0.2 255.255.255.0
 no ip redirects
!
interface Vlan110
 ip address 10.110.0.1 255.255.255.0
!
tftp-server disk0:WWW.P12
tftp-server disk0:WWWIN.P12
alias exec csm5 show module csm 5
alias exec csmrun show run | begin module ContentSwitchingModule
!
```

Aggregation2

```
agg2#show run
!
ssl-proxy module 3 allowed-vlan 3,110
vtp domain mydomain
vtp mode transparent
ip subnet-zero
!
!
no ip domain-lookup
!
mls flow ip destination
mls flow ipx destination
mls sampling packet-based 1024 4096
!
spanning-tree mode rapid-pvst
spanning-tree loopguard default
spanning-tree vlan 3,5,6,10,12,80,100 priority 16384
!
module ContentSwitchingModule 5
 vlan 5 client
  ip address 10.20.5.5 255.255.255.0
  alias 10.20.5.6 255.255.255.0
!
vlan 3 server
  ip address 10.20.3.5 255.255.255.0
  alias 10.20.3.6 255.255.255.0
!
vlan 6 client
  ip address 10.20.6.5 255.255.255.0
  alias 10.20.6.6 255.255.255.0
!
vlan 10 server
  ip address 10.20.5.5 255.255.255.0
!
vlan 12 server
  ip address 10.20.6.5 255.255.255.0
!
```

```
 probe TCP tcp
  interval 2
  failed 3
!
 probe ICMP icmp
  interval 3
  failed 5
!
serverfarm SSLBLADE-VIP1
  nat server
  no nat client
  real 10.20.3.80 445
   inservice
  real 10.20.3.90 445
   inservice
  probe TCP
!
serverfarm SSLBLADE-VIP2
  nat server
  no nat client
  real 10.20.3.80 446
   inservice
  real 10.20.3.90 446
   inservice
  probe TCP
!
 serverfarm WEB-VIP1
  nat server
  no nat client
  real 10.20.5.14
   inservice
  real 10.20.5.15
   inservice
  probe TCP
!
 serverfarm WEB-VIP2
  nat server
no nat client
  real 10.20.5.16
   inservice
  real 10.20.5.17
   inservice
  probe TCP
!
 sticky 1 cookie cookie-server timeout 10
 sticky 2 cookie cookie-server timeout 10
!
vserver HTTP-VIP1
  virtual 10.20.5.80 tcp www
  vlan 5
  serverfarm WEB-VIP1
  advertise active
  sticky 10 group 1
  replicate csrp sticky
  persistent rebalance
  parse-length 4000
  inservice
!
vserver HTTP-VIP2
  virtual 10.20.5.90 tcp www
  vlan 5
  serverfarm WEB-VIP2
  advertise active
  sticky 10 group 2
```

```
  replicate csrp sticky
  persistent rebalance
  parse-length 4000
  inservice
!
vserver HTTPS-VIP1
  virtual 10.20.5.80 tcp https
  vlan 5
  serverfarm SSLBLADE-VIP1
persistent rebalance
  inservice
!
vserver HTTPS-VIP2
  virtual 10.20.5.90 tcp https
  vlan 5
serverfarm SSLBLADE-VIP2
persistent rebalance
  inservice
!
vserver WEBDECRYPT-VIP1
  virtual 10.20.3.81 tcp 445
  vlan 3
  serverfarm WEB-VIP1
  sticky 10 group 1
  replicate csrp sticky
  persistent rebalance
  parse-length 4000
  inservice
!
vserver WEBDECRYPT-VIP2
  virtual 10.20.3.81 tcp 446
  vlan 3
serverfarm WEB-VIP2
  sticky 10 group 2
  replicate csrp sticky
  persistent rebalance
  parse-length 4000
  inservice
!
ft group 1 vlan 100
  priority 10
  heartbeat-time 5
  failover 4
!
redundancy
 mode rpr-plus
 main-cpu
  auto-sync startup-config
  auto-sync running-config
  auto-sync standard
!
vlan dot1q tag native
!
vlan 3
 name SSLVLAN
!
vlan 5
 name WEBTIERVLAN
!
vlan 6
 name APPTIERVLAN
!
vlan 10
 name WEBSERVERVLAN
```

```
!
vlan 12
 name APPSERVERVLAN
!
vlan 80
 name Layer3_VLAN
!
vlan 100
 name CSM_fault_tolerant
!
vlan 110
 name ssl_admin
!
interface Port-channel2
 no ip address
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 3,5,6,10,12,80,100,110
 switchport mode trunk
 spanning-tree guard loop
!
interface GigabitEthernet1/1
 description to_mp_agg2
 no ip address
 logging event link-status
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 3,5,6,10,12,80,100
 switchport mode trunk
 channel-group 2 mode active
 channel-protocol lacp
!
interface GigabitEthernet4/1
 description to_mp_acc1
 no ip address
 switchport
!
interface GigabitEthernet4/6
 description to_mp_agg2
 no ip address
 logging event link-status
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk allowed vlan 3,5,6,10,12,80,100,110
 switchport mode trunk
 channel-group 2 mode active
 channel-protocol lacp
!
interface Vlan5
 ip address 10.20.5.3 255.255.255.0
 no ip redirects
 standby 1 ip 10.20.5.1
 standby 1 priority 100
 standby 1 preempt
!
interface Vlan6
 ip address 10.20.6.3 255.255.255.0
 no ip redirects
 standby 1 ip 10.20.6.1
 standby 1 priority 100
 standby 1 preempt
!
interface Vlan80
 ip address 10.80.0.3 255.255.255.0
```

```
 no ip redirects
!
interface Vlan110
 ip address 10.110.0.2 255.255.255.0
!
tftp-server disk0:WWW.P12
tftp-server disk0:WWWIN.P12
alias exec csm5 show module csm 5
alias exec csmrun show run | begin module ContentSwitchingModule
!
```

SSL Offloader 1

```
ssl-off1#show run br
!
hostname ssl-off1
!
ssl-proxy service VIP10.20.5.80
 virtual ipaddr 10.20.3.80 protocol tcp port 445
 server ipaddr 10.20.3.81 protocol tcp port 445
 certificate rsa general-purpose trustpoint www
 inservice
!
ssl-proxy service VIP10.20.5.90
 virtual ipaddr 10.20.3.80 protocol tcp port 446
 server ipaddr 10.20.3.81 protocol tcp port 446
 certificate rsa general-purpose trustpoint wwwin
 inservice
!
ssl-proxy vlan 3
 ipaddr 10.20.3.8 255.255.255.0
 gateway 10.20.3.6
!
ssl-proxy vlan 110
 ipaddr 10.110.0.8 255.255.255.0
 gateway 10.110.0.1
 admin
!
crypto ca trustpoint www
 rsakeypair www
!
crypto ca trustpoint wwwin
 rsakeypair wwwin
!
crypto ca certificate chain www
 certificate 02
 certificate ca 00
!
crypto ca certificate chain wwwin
 certificate 03
 certificate ca 00
```

SSL Offloader 2

```
ssl-off2#show run br
!
hostname ssl-off2
!
ssl-proxy service VIP10.20.5.80
```

```
 virtual ipaddr 10.20.3.90 protocol tcp port 445
 server ipaddr 10.20.3.81 protocol tcp port 445
 certificate rsa general-purpose trustpoint www
 inservice
!
ssl-proxy service VIP10.20.5.90
 virtual ipaddr 10.20.3.80 protocol tcp port 446
 server ipaddr 10.20.3.81 protocol tcp port 446
 certificate rsa general-purpose trustpoint wwwin
 inservice
!
ssl-proxy vlan 3
 ipaddr 10.20.3.9 255.255.255.0
 gateway 10.20.3.6
!
ssl-proxy vlan 110
 ipaddr 10.110.0.9 255.255.255.0
 admin
!
crypto ca trustpoint www
 rsakeypair www
!
crypto ca trustpoint wwwin
 rsakeypair wwwin
!
crypto ca certificate chain www
 certificate 02
 certificate ca 00
!
crypto ca certificate chain wwwin
 certificate 03
 certificate ca 00
```

Numerics

7600 **3-1**
802.1q **1-11**
802.1q tag **3-4, 3-5**
802.1s **1-11**
802.1w **1-11**
802.3ad **1-11**

A

AAA **1-13**
ABR **2-4, 2-11**
access control list **3-13**
access layer **1-7**
accounting management **1-14**
ACL **3-13**
ACLs **1-13**
active-standby **3-3**
aggregation layer **1-6**
Apache **4-6**
application layer **1-7, 1-8**
Application Optimization Services **1-10**
application requirements **1-3**
area border router **2-4**
ARP test **2-7**
ASBR **2-3**
asynchronous communications **1-9**
autonomous system border routers **2-3**
availability **1-3**

B

back-end layer **1-8**
BGP **1-11**
BGP session **2-4**
Border Gateway Protocol **1-11**
border routers **2-3**
bridge mode **3-3, 3-6, 4-4**
broadcast domains **4-11**
broadcast ping test **2-7**
Broadcast Suppression **1-11**
building blocks **1-2**
business continuance **1-9**
business logic **1-7**
business requirements **1-3**

C

caching **1-11**
Call Managers **1-7**
campus-to-campus connectivity **1-9**
Catalyst 6500 **3-1, 3-3**
CatOS **3-1**
CC **3-3, 3-4**
CDP **1-14**
certificates **4-1**
Clear text traffic **4-2**
client NAT **3-10**
Coarse Wave Division Multiplexing **1-10**
concurrent connections **3-3**
configuration management **1-14**
connections per second **3-3**
Content Distribution Managers **1-7**

content engines **1-6, 3-5**
Content Services Switch **2-1**
content switches **1-6**
content switching **1-11, 3-1**
Content Switching Module **2-1, 3-1, 4-1**
content transformation devices **3-5**
core switches **1-3**
CPS **3-3, 3-4**
CSM **2-1, 2-5, 3-1, 4-1**
CSS **2-1**
CWDM **1-10**

D

data encryption **4-1**
decrypted traffic **4-13**
demilitarized zone **2-1**
Denial of Service **1-13**
Dense Wave Division Multiplexing **1-10**
designated router **2-8**
disaster recovery **1-9**
distributed data centers **1-9**
distribution layer **1-6**
DMZ **2-1, 2-2, 2-5, 2-12**
DWDM **1-10**
dynamic routing **2-1, 2-12**

E

EIGRP **1-11, 2-4**
equal cost path load balancing **2-4**
ESCON **1-10**
Extranet Server Farm **1-3**

F

failover tests **2-7**
fault management **1-14**
fault tolerant VLAN **3-11**
FCAPS **1-14**
federal government agencies **1-1**
Fibre Channel **1-12**
Fibre-Channel (FC) **1-8**
financial institutions **1-1**
Firewalls **1-13**
firewalls **1-6, 1-7**
Firewall Services Module **2-1, 3-14**
FlexWAN card **2-4**
front-end layer **1-7**
FTP **1-7**
FWSM **2-1, 2-2, 2-7, 2-14, 3-14, 4-1**
FWSMs **2-6**

H

handshake **4-1**
healthcare **1-1**
hello packet **2-6**
high availability **1-1, 3-2, 3-3**
high speed connection **1-9**
Host IDS **1-13**
hot standby router protocol **3-6**
HSRP **1-11, 3-6, 3-8**
HTTP **4-2**
HTTP daemon **4-6**
HTTP GET **4-2**
HTTP responses **4-14**
HTTPS **4-2, 4-8**
HTTP session persistence **4-1**
HTTPS requests **4-12**

I

ICMP ping **4-5**
IDS **1-7, 1-13**
IDSs **1-6**

IGP **1-11**
Infrastructure **1-10**
Intelligent Network Services **1-10**
Interior Gateway Protocols **1-11**
Internet Server Farm **1-3**
Intranet Server Farm **1-3**
IPSEC **1-14**
IP spoofing **1-13**
IPTV Broadcast servers **1-7**
iSCSI **1-8**
IS-IS **2-4**

K

key initiation **4-1**

L

Layer 2 **1-7, 1-10**
Layer 2 attacks **1-13**
Layer 3 **1-10, 1-11**
Layer 3 VLAN **2-1**
Layer 4 **1-11**
Layer 5 **1-7, 1-11**
Link Aggregate Control Protocol **1-11**
LoopGuard **1-11**
low convergence **3-2**

M

mainframes **1-8**
Management **1-10**
management **1-1, 1-3**
Management services **1-14**
medium sized servers **1-8**
message integrity **4-1**
Metro **1-10**
metro optical **1-9**
metro transport layer **1-9**
MHSRP **1-11, 2-5**
middleware **1-7**
monitoring **1-14**
MPEG **1-12**
MSFC **2-1, 2-2, 3-1, 4-4, 4-7**
MSFC-inside **2-2, 2-4, 2-5, 2-12**
MSFC-outside **2-2, 2-4, 2-5**
Multicast **1-7**
multigroup hot standby router protocol **2-5**
Multilayer Switching Feature Card **2-1**

N

NAS **1-12**
NAT **2-10**
Native IOS **3-1**
network address translation **2-10**
Network Attached Storage **1-12**
network reconnaissance **1-13**
nssa **2-11**

O

one-arm mode **3-3**
One Time Passwords **1-14**
OSM card **2-4**
OSPF **1-11, 2-1, 2-4, 2-12, 2-14, 2-28**

P

packets per second **3-3**
PAT **3-10**
PBR **1-11**
performance management **1-14**
PIX **2-1**
policy based routing **1-11**
port 446 **4-6**

port 80 **4-2**
port address translation **3-10**
PPS **3-3, 3-4**
Private VLANs **1-11**
Private WAN **1-3**

Q

QoS **1-7, 1-11, 3-12**
QoS ACLs **1-13**
Quicktime **1-12**

R

relational database **1-8**
Remote Access **1-3**
Renewable Energy Policy Project **1-1**
replication **1-9**
revenue **1-1**
reverse proxy caching **1-12**
RFC 1918 **3-10**
routed mode **4-3**
Router ACLs **1-13**
RSA private key **4-1**

S

SAN **1-12**
scalability **1-1, 1-3**
secure router mode **3-3, 3-7**
secure socket layer **4-1**
Security **1-10**
security **1-1, 1-3, 1-7**
security domain **2-1, 2-5**
security levels **2-3**
security management **1-14**
Security services **1-12**
security templates **1-14**
server authentication **4-1**
server default gateway **3-3**
Server load balancing **1-11**
server NAT **3-10**
server-to-server communication **3-8, 4-4, 4-9**
service provisioning **1-14**
SMTP **1-7**
SONET **1-10**
spanning-tree **4-9**
SSH **1-14**
SSL **1-7, 4-1**
SSLBLADE-VIP1 **4-13**
SSL decryption **4-6**
SSL encryption engine **4-1**
SSL offloaders **1-6, 3-5**
SSL offloading **1-11**
ssl-proxy vlan **4-11**
SSL Service Module **4-1**
SSLSM **4-1**
SSLSM redundancy **4-5**
stateful failover **3-3**
static routes **2-9**
static routing **1-11, 2-12**
Storage **1-10**
Storage Area Networks **1-12**
storage layer **1-8**
storage mirroring **1-9**
Storage services **1-12**
storage-to-storage **1-12**
stub area **2-11**
SVI **2-1, 2-8, 4-1**
switched VLAN interface **2-1, 4-1**
synchronous communications **1-9**

T

TCP handshakes **4-5**
TCP probes **4-5**
Telnet **1-7**

TN3270 **1-7**
totally stubby area **2-11**
total number of VIPs **3-3**
troubleshooting **1-14**

U

UDLD **1-11**
Unauthorized access **1-13**
Uni-Directional Link Detection **1-11**
user-to-server **1-12**
utilities **1-1**

V

VIP **3-10**
Virtual IP address **2-5, 2-14**
viruses and worm **1-13**
VLAN ACLs **1-13**
VLANs **1-7**
VRRP **1-11**
Vserver IP address **3-10**
VTY security **1-14**

W

Web servers **1-7**
Windows Media **1-12**